



ULISSE s.r.l.
Servizio Distribuzione Apple
Via Pacini, 22
20131 Milano
Tel. (02) 2367783-2367406

su licenza esclusiva
Apple Computer S.p.A.

Inside AppleTalk

IAT01

Inside AppleTalk

Preliminary Notes
July 14, 1986

Writers: Gursharan S. Sidhu
Richard F. Andrews
Alan B. Oppenheimer

Copyright © 1986 by Apple Computer, Inc.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without prior written permission of Apple Computer, Inc.

Apple, AppleTalk, the Apple logo, and LaserWriter are trademarks of Apple Computer, Inc.

Macintosh is a trademark of McIntosh Laboratory, Inc., and is being used with express permission of its owner.

PostScript is a trademark of Adobe Systems, Inc.

The authors wish to thank Gene Tyacke for writing the chapter on the AppleTalk Poke program and for updating the description of the AppleTalk Peek utility. Bruce Gaya's contributions to the experimental validation of RTMP, ZIP, and NBP operation in the case of half-bridges are gratefully acknowledged. Also we wish to thank Bob Anders for his editorial contributions.

This book was produced using the Apple Macintosh computer, AppleTalk, and the LaserWriter printer.

Contents

vii Preface

I-1	I. AppleTalk - An Introduction
I-1	Hardware Specifications
I-1	AppleTalk Protocol Architecture
I-2	Physical Layer
I-2	AppleTalk Link Access Protocol (ALAP)
I-2	Datagram Delivery Protocol (DDP)
I-3	Routing Table Maintenance Protocol (RTMP)
I-3	Name-Binding Protocol (NBP)
I-3	AppleTalk Transaction Protocol (ATP)
I-3	Zone Information Protocol (ZIP)
I-3	Echo Protocol (EP)
I-3	AppleTalk Data Stream Protocol (ADSP)
I-4	Additional Protocols
II-1	II. Electrical Specifications
II-1	Bit Encoding and Decoding
II-1	Signal Transmission and Reception
II-1	Carrier Sense
III-1	III. AppleTalk Link Access Protocol (ALAP)
III-1	About AppleTalk Link Access Protocol
III-1	Bus Access Management
III-1	Node Addressing Mechanism
III-1	Dynamic Node ID Assignment
III-2	Data Transmission and Reception
III-3	ALAP Frame Format
III-3	ALAP Header
III-4	Frame Check Sequence Field
III-4	Frame Size Limitations
III-4	Frame Transmission
III-4	Directed Data Frames
III-5	Synchronization
III-5	Directed Transmissions
III-6	Broadcast Transmissions
III-6	Frame Reception
III-6	Frame Size Error
III-6	Overflow/Underflow Error
III-6	Frame Type
III-7	Frame Check Sequence
III-7	Algorithms
III-7	CRC-CCITT Calculation
III-8	Random Wait Time Determination
III-9	ALAP Procedural Model
III-9	Assumptions
III-9	Global Constants, Types and Variables
III-10	Hardware Interface Declarations
III-11	Interface Procedures and Functions
III-12	InitializeLAP Procedure

III-12	AcquireAddress Procedure
III-13	TransmitPacket Function
III-13	TransmitLinkMgmt Function
III-16	TransmitFrame Procedure
III-17	ReceivePacket Procedure
III-17	ReceiveLinkMgmt Function
III-18	ReceiveFrame Function
III-20	Miscellaneous Functions
III-21	SCC Implementation
IV-1	IV. Datagram Delivery Protocol (DDP)
IV-1	About Datagram Delivery Protocol
IV-1	Sockets and Socket Identification
IV-2	DDP Protocol Type Field
IV-2	Socket Listeners
IV-2	DDP Interface
IV-2	Opening a Statically-Assigned Socket
IV-3	Opening a Dynamically-Assigned Socket
IV-3	Closing a Socket
IV-3	Sending a Datagram
IV-3	Datagram Reception by the Socket Listener
IV-3	DDP Internal Algorithm
IV-4	DDP Packet Format
IV-4	Short and Long Form Headers
IV-5	DDP Checksum Computation
IV-5	Hop Counts
IV-5	DDP Routing Algorithm
IV-7	Use of Name-Binding Protocol in Conjunction with Sockets
IV-7	Network Number Equivalence
V-1	V. Routing Table Maintenance Protocol (RTMP)
V-1	About Routing Table Maintenance Protocol
V-1	Bridges
V-1	Local Bridges
V-1	Half Bridges
V-1	Backbone Bridges
V-2	Bridge Model
V-2	Internet Topologies
V-2	Routing Tables
V-3	Routing Table Maintenance
V-4	"Aging" Routing Table Entries
V-4	Values for the Validity and Send-RTMP Timers
V-4	RTMP Data Packet Format
V-5	Sender's Network Number
V-5	Sender's Node ID
V-5	Routing Tuples
V-5	Assignment of Network Numbers
V-6	RTMP Initialization and Maintenance Algorithms
V-6	Initialization
V-6	Maintenance
V-7	Update-the-Entry
V-7	Create-New-Entry
V-7	Replace-Entry
V-8	Supplements to the Algorithm for Half-Bridges

V-9	Additional RTMP Services
V-9	RTMP and Non-Bridge Nodes
V-10	RTMP Routing Algorithm
VI-1	VI. Name Binding Protocol (NBP)
VI-1	About Name Binding Protocol
VI-1	Network-Visible Entities
VI-1	Entity Names
VI-2	Name Binding
VI-2	Names Directory and Names Tables
VI-3	Aliases and Enumerators
VI-3	Names Information Socket
VI-3	Name Binding Services
VI-3	Name Registration
VI-3	Name Deletion
VI-4	Name Lookup
VI-4	Name Confirmation
VI-4	NBP on a Single Network
VI-5	NBP on an Internet
VI-5	Zones
VI-5	Name Lookup on an Internet
VI-6	NBP Interface
VI-6	Registering a Name
VI-6	Removing a Name
VI-7	Name Lookup
VI-7	Confirming a Name
VI-7	NBP Packet Formats
VI-8	Control
VI-8	Tuple Count
VI-8	NBP ID
VI-8	NBP Tuple
VII-1	VII. AppleTalk Transaction Protocol (ATP)
VII-1	About AppleTalk Transaction Protocol
VII-1	Transactions
VII-1	At-Least-Once Transactions (ALO)
VII-2	Exactly-Once Transactions (XO)
VII-3	Multi-Packet Response
VII-3	Transaction IDs
VII-4	ATP BitMap/Sequence Number
VII-5	Responders with Limited Buffer Space
VII-5	ATP Packet Format
VII-6	ATP Interface
VII-7	Send a Request
VII-8	Open a Responding Socket
VII-8	Close a Responding Socket
VII-8	Receive a Request
VII-9	Send a Response
VII-9	ATP State Model
VII-10	ATP Requester
VII-11	ATP Responder
VII-13	Some Optional ATP Interface Calls
VII-13	Release a RspCB
VII-13	Release a TCB

VII-13	Wrap Around and Generation of Transaction Identifiers
VIII-1	VIII. Zone Information Protocol (ZIP)
VIII-1	About Zone Information Protocol
VIII-1	ZIP Services
VIII-1	The Network-Number to Zone-Name Mapping
VIII-1	The Zone Information Table
VIII-1	The Zone Information Socket: ZIP Queries and Replies
VIII-2	ZIT Maintenance
VIII-2	Changing Zone Names
VIII-4	Listing Zone Names
VIII-5	Packet Format
VIII-5	NBP Routing in Bridges
VIII-6	Zone Name Assignment
VIII-6	Timer Values
VIII-6	Implementation Notes
IX-1	IX. Printer Access Protocol (PAP)
IX-1	About Printer Access Protocol
IX-2	The Protocol
IX-3	Connection Establishment (Opening) Phase
IX-4	Data Transfer Phase
IX-5	Duplicate Filtration
IX-6	Connection Termination (Closing) Phase
IX-6	Status Gathering
IX-7	PAP Packet formats
IX-7	PAP Type Values
IX-7	The Client-PAP Interface
IX-7	PAPOpen
IX-8	PAPClose
IX-8	PAPRead
IX-8	PAPWrite
IX-9	PAPStatus
IX-9	SLInit
IX-10	GetNextJob
IX-10	SLClose
IX-10	PAPRegName
IX-10	PAPRemName
IX-11	HeresStatus
IX-11	The Apple LaserWriter™
X-1	X. Echo Protocol (EP)
X-1	About Echo Protocol
XI-1	XI. AppleTalk Session Protocol (ASP)
XI-1	About AppleTalk Session Protocol
XI-1	Sessions, Commands and Command Replies
XI-2	What ASP Does Not Do
XI-2	ASP Services and Features
XI-3	Opening and Closing Sessions
XI-3	Session Maintenance
XI-4	Session Requests on an Open Session
XI-4	ASP Commands
XI-4	ASP Writes

XI-5	Sequencing and Duplicate Filtration on Sessions
XI-5	Getting Service Status Information
XI-5	ASP Client Interface
XI-5	Server Side
XI-6	SPGetParms
XI-6	SPInit
XI-7	SPGetSession
XI-7	SPCloseSession
XI-8	SPGetRequest
XI-8	SPCmdReply
XI-9	SPWrtContinue
XI-10	SPWrtReply
XI-10	SPNewStatus
XI-11	SPAttention
XI-11	Workstation Side
XI-11	SPGetParms
XI-11	SPGetStatus
XI-12	SPOpenSession
XI-12	SPCloseSession
XI-13	SPCommand
XI-14	SPWrite
XI-14	Packet Formats and ASP Internals
XI-14	Opening a Session
XI-15	Getting Service Status
XI-16	ASP Commands
XI-16	ASP Writes
XI-17	Reply Size Error Checking
XI-17	Tickles and Session Maintenance
XI-18	Attention Requests
XI-18	Closing a Session
XI-19	SPCmdType Values
XI-19	SPErrors Values
XI-20	Time-Outs and Retry Counts

XII-1 XII. AppleTalk Data Stream Protocol (ADSP)

XIII-1 XIII. AppleTalk Filing Protocol (AFP)

A-1 Appendix A:

B-1 Appendix B: Miscellaneous AppleTalk Parameters

C-1 Appendix C: AppleTalk Peek

D-1 Appendix D: AppleTalk Poke

Preface

AppleTalk has been under development for over two years starting late in the Fall of 1983. *Inside AppleTalk* was first published in May of 1984 and made available at that time to potential developers of AppleTalk products. Since that date, the protocols have been implemented and tested with care by Apple Computer and by third party developers. This real-life experience has led to several enhancements and refinements of the protocols. It has, furthermore, revealed certain inaccuracies in the documentation. Several additional protocols have been added to the architecture.

For these reasons *Inside AppleTalk* has been under constant revision and correction for several months. Although a more carefully edited and revised version will eventually be published by Apple at a later date, we consider it important to get the information contained herein into the hands of developers as soon as possible. We are pleased to make this early form of the revised manual available to developers at this time. Please communicate any errors, omissions, unclear descriptions, etc., to the authors at the following address (no phone calls please):

Gursharan S. Sidhu,
Apple Computer Inc., M/S 22-O,
20525 Mariani Avenue,
Cupertino, CA. 95014.

Use this address only to send us comments on this document. Please do *not* write to this address or call us for technical consultations and questions regarding AppleTalk. Those issues should be addressed to Apple's Technical Support Group.

This document has been divided into two parts. The first provides detailed specifications of the various protocols constituting the AppleTalk protocol architecture. This is the primary source for specification and description of the protocols themselves. We have strived hard to be machine independent. Although occasional hints are provided on implementation issues, this is not a description of how these protocols are actually implemented on any computer system.

The second part contains the users manuals of the two main AppleTalk development tools: Peek and Poke. These tools were first provided by us in the Spring of 1984 and have undergone many modifications and enhancements since that time. The descriptions provided here refer to the latest available versions (version ?? for Peek, and version ?? for Poke).

Although the changes to *Inside AppleTalk* are too numerous to list in detail here, here is a brief discussion of the principal ones.

1. *ALAP Clarifications*

The description of this protocol in our original 1984 document was imprecise in some respects. This chapter has been carefully rewritten and the protocol details specified in a detailed algorithmic manner using a Pascal-like language. The role of the synch pulse has been discussed more carefully.

2. *New Protocols*

Five new protocols have been added to AppleTalk and one protocol has been replaced. These are:

- PAP -- the Printer Access Protocol
- ZIP -- the Zone Information Protocol
- ASP -- the AppleTalk Session Protocol
- AFP -- AppleTalk Filing Protocol
- EP -- the Echo Protocol
- ADSP -- the AppleTalk Data Stream Protocol

Complete specifications of these are included here. Of these five protocols, PAP and ZIP has been fully implemented and tested. ASP and AFP are currently being implemented and tested. The EP and ADSP chapters should at this stage be considered as proposals. Apple plans to implement and test them in the near future. The discussion of these two protocols is provided here for the purpose of information and to invite your comments on EP and ADSP.

It should also be noted that ADSP replaces our early document on DSP (Data Stream Protocol), which was never implemented and is no longer a part of the AppleTalk protocol architecture. Readers are advised to discard the old chapter on DSP to avoid future confusion with the newly defined ADSP.

3. *Internetting Issues*

Interconnection of several AppleTalk networks into a larger internet system has been one of the areas of most intense AppleTalk activity at Apple. Several changes and refinements have been made to DDP, NBP and RTMP to correctly handle certain issues related to the use of network numbers and zone names. The Zone Information Protocol has been completely designed, implemented and tested. It is important for developers to carefully review the relevant chapters and ensure that their implementations of these protocols conform to the specifications in this document. Failure to do so will, in all likelihood, lead to improper operation of their devices on AppleTalk internets. With AppleTalk internet routers and back-bone bridges appearing in the market place at this time these issues are expected to rapidly assume major significance.

In the discussion of DDP, the issue of when two network numbers are equivalent, and how to use network numbers when sending or receiving has been specified in detail.

Although RTMP has remained essentially the same as in the May 1984 version of *Inside AppleTalk*, several issues of detail have been experimentally examined and specified here. More specifically:

- A bridge's routing algorithm has been more clearly stated
- All algorithms have been better defined (especially as related to non-seed bridges)
- The aging of ABridge is now required in non-bridge nodes
- A new mechanism has been added for non-bridge nodes to rapidly acquire their network number
- A discussion of internet topologies is provided
- Mechanisms for the support for ZIP commands are discussed.

With respect to NBP, the following issues have been elucidated:

- The use of zone * in NBP calls
- The algorithm used by bridges for converting NBP BrRq's to zone-wide NBP LkUp's
- NBP implementations should not respond to their own lookup requests
- Case insensitivity of entity names
- NBP confirmation returns socket number
- Null zone field same as *

4. ATP Clarifications

Several subtle issues related to ATP have been examined and clarified herein. These include a detailed discussion of transaction ID wrap around and an algorithm for TID generation that eliminates wrap around problems. Other details of interest are:

- ATP calls to DDP should ignore errors
- ATP does not return an error if exactly-once mode is requested but not available
- STS bit usage
- EOM bit is not available for use by ATP clients
- Infinite retries and optional socket on Send-Request call.

5. Printer Access Protocol (PAP)

The specification of the printer access protocol had formerly been included in the Inside LaserWriter manual. However, this protocol is now used to communicate over AppleTalk with other devices such as the ImageWriter-II. Also, the earlier discussion of PAP mixed the protocol's specification with a description of the somewhat restricted implementation of the protocol on the Macintosh. This caused the false impression that PAP allowed only one connection at a particular time. For these reasons, this chapter has been completely rewritten, making it implementation independent and more complete.

6. Table of Different Reserved Fields, Sockets, Protocol Types

An appendix listing the different reserved values of statically-assigned sockets, ALAP protocol types, and DDP protocol types is included in the Appendix.

I. AppleTalk - An Introduction

The AppleTalk Personal Network has been designed as a simple, inexpensive, and flexible way to interconnect computers, peripheral devices, and servers. AppleTalk can be used in three major configurations:

- As a stand-alone local-area network, AppleTalk allows workgroups to exchange information and share resources like printers, file servers, modems, and other peripherals. In this configuration, a stand-alone AppleTalk network can be used to interconnect up to 32 nodes (computers, devices, and servers).
- Through the use of bridging devices, large numbers of AppleTalk networks, each with up to 32 nodes, can be connected into large, complex internetworks. The nodes connected to this internetwork obtain the same services as on a stand-alone AppleTalk. Through gateway devices, AppleTalks can be hooked into other local-area networks such as Ethernet™.
- As a peripheral bus, AppleTalk can connect a Macintosh to its dedicated attached devices.

To achieve this range of flexibility, Apple has designed highly reliable, inexpensive hardware, and a sophisticated set of communications protocols.

Hardware Specifications

At the physical level, AppleTalk has a bus topology consisting of a linear trunk cable with intervening connection boxes to which devices attach via a short drop cable. Since end-user installation is assumed, assembled trunk cables are sold in standard 2, and 10 meter lengths with molded miniature DIN connectors. The trunk cables consist of shielded, twisted-pair cable. Devices are connected to the trunk cable with AppleTalk connection boxes. The connection box is a small plastic case containing a transformer, resistive and capacitive circuits for noise immunity, and two 3-pin miniature DIN connectors with terminating switches to a 100 ohm terminating resistor. Attached to this box is an 18-inch drop cable which terminates at the device with either a DB-9, DB-25 or miniature DIN connector plug. Rolls of trunk cable and miniature DIN connectors are available allowing users to build custom network cable runs of lengths other than the standard 2 and 10 meters.

For a complete hardware specification, refer to the *AppleTalk Electrical/Mechanical Specification* and the *AppleTalk Transformer Specification*, included as Appendixes B and C at the end of this manual.

AppleTalk Protocol Architecture

Underlying all use of AppleTalk is a specific set of rules, or communication protocols. Apple has developed a set of protocols that correspond to the various layers of the International Standards Organization (ISO) Open Systems Interconnection (OSI) reference model. Protocols at the ISO-OSI layers 1 through 5 (Physical, Data Link, Network, Transport, and Session) form the core of the AppleTalk protocol architecture.

While Apple recommends the use of these protocols, communication over AppleTalk is not dependent on their exclusive use. The AppleTalk architecture is designed to allow developers to add special functions and features for their particular applications. All protocols are layered, functionally distinct entities, allowing easy access to and addition of alternative protocols. A software developer could, for instance, leverage upon the physical and data link layers to build a different protocol architecture.

When designing alternative protocols within a layered architecture, it is important to avoid implementing layers as separate modules with interfaces based on high-overhead calls to the operating system. A layered architecture is modular only with respect to the clear and distinct functionality provided at each layer. The intelligent design of interfaces between layers can avoid problems such as unnecessary *buffer copying* when passing data between layers.

The general method used by the AppleTalk architecture for moving user information up or down through the protocol layers is known as data encapsulation/decapsulation. A unit of user information is enclosed by a layer-specific header and/or trailer as it moves through each layer, from the user application down to the Data Link layer. The corresponding protocol layers at the receiving end examine and remove the layer-specific protocol information, as the user data moves up through the layers to the receiving user application.

The protocols comprising the AppleTalk architecture are outlined below. Figure I-1 summarizes this architecture.

Physical Layer

The electrical and mechanical characteristics of AppleTalk correspond to the Physical layer of the ISO-OSI reference model. The Physical layer performs the functions of bit encoding/decoding, synchronization, signal transmission/reception, and carrier sense. The layered nature of the AppleTalk architecture allows for the Physical layer to be replaced by another medium as long as these functions are provided, and the interface between the Physical and Data Link layers is maintained.

A detailed discussion of the functions performed by the Physical layer is provided in the Electrical Specifications chapter.

AppleTalk Link Access Protocol (ALAP)

The AppleTalk Link Access Protocol (ALAP) corresponds to the Data Link layer of the ISO-OSI reference model. ALAP, which must be common to all systems on the bus, provides *best effort* delivery of information between devices, also known as *nodes*. It provides the basic underlying service of packet transmission between the devices connected to a single AppleTalk network. ALAP manages the encapsulation of data in a *frame* and then provides access to the bus for transmission and reception of frames.

Datagram Delivery Protocol (DDP)

The Datagram Delivery Protocol (DDP) is found at the next level of the architecture, corresponding to the Network layer of the ISO-OSI model. While the ALAP protocol provides delivery of packets over a single AppleTalk network, the Datagram Delivery Protocol extends this mechanism to include a group of interconnected AppleTalk networks,

known as an *internet*. DDP also introduces the notion of logical addressable entities within nodes, known as *sockets*. *Datagrams* are packets of data exchanged between sockets; sockets may therefore be visualized simply as the source and destination of datagrams.

Routing Table Maintenance Protocol (RTMP)

The AppleTalk architecture includes several protocols at the level above DDP, corresponding to the Transport layer of the ISO-OSI model. These protocols, add different levels or types of functionality to the underlying datagram delivery service. This Transport layer protocol allows bridges/internet routers to dynamically discover routes to the different AppleTalk networks in an internet. Non-bridge nodes use a subset of RTMP (known as the *RTMP stub*) to determine the number of the network to which they are connected, as well as the node ID of a bridge on their network.

Name-Binding Protocol (NBP)

Name-Binding Protocol (NBP), a Session layer protocol, lets network users use character string names for socket clients and network services. The basic function of NBP is the translation of a character string name into the internet socket address of the corresponding socket client.

NBP also introduces the concept of a *zone* - an arbitrary subset of networks in an internet where each network is in one and only one zone.

AppleTalk Transaction Protocol (ATP)

The AppleTalk Transaction Protocol (ATP) adds a measure of reliability by providing a loss-free *transaction* service between sockets. This allows exchanges between two socket *clients*, in which one client requests the other to perform a particular task and report the result.

Zone Information Protocol (ZIP)

This protocol maintains an internet-wide mapping of networks to zone names. ZIP is used by Name-Binding Protocol to determine which networks belong to a given zone.

Echo Protocol (EP)

This is a very simple protocol that allows any node to send a datagram to any other node on an AppleTalk internet and receive an *echoed* copy of that packet in return. It is useful for probing to confirm the existence of a particular node, or to make round trip delay measurements.

AppleTalk Data Stream Protocol (ADSP)

This is a connection-oriented protocol providing a reliable, full-duplex, byte stream service between any two sockets on an AppleTalk internet. It ensures in-sequence, duplicate-free delivery of data over its connections.

Additional Protocols

The above mentioned protocols constitute the *core* AppleTalk protocols, upon which most higher-level services will be built. In addition, Apple has developed, or is developing various of these higher-level protocols, as summarized below:

Printer Access Protocol (PAP) - utilizes ATP XO to create a stream-like service for talking with the Apple LaserWriter™ and other stream-based devices.

PostScript™ - developed by Adobe Systems® of Palo Alto. Presents a resolution-independent standard method of describing graphical and textual data. Utilized in talking to the Apple LaserWriter™, and other graphical devices. Detailed in *Inside LaserWriter*, published by Apple Computer and the *PostScript Language Reference Manual*, published by Addison-Wesley.

AppleTalk Session Protocol (ASP) - a general protocol, built on ATP, providing session establishment, maintenance and teardown, along with request sequencing.

AppleTalk Filing Protocol (AFP) - a presentation level protocol for accessing remote file systems, built on ASP.

Detailed discussions of each of the AppleTalk protocol layers, with the exception of PostScript, are provided in subsequent chapters.

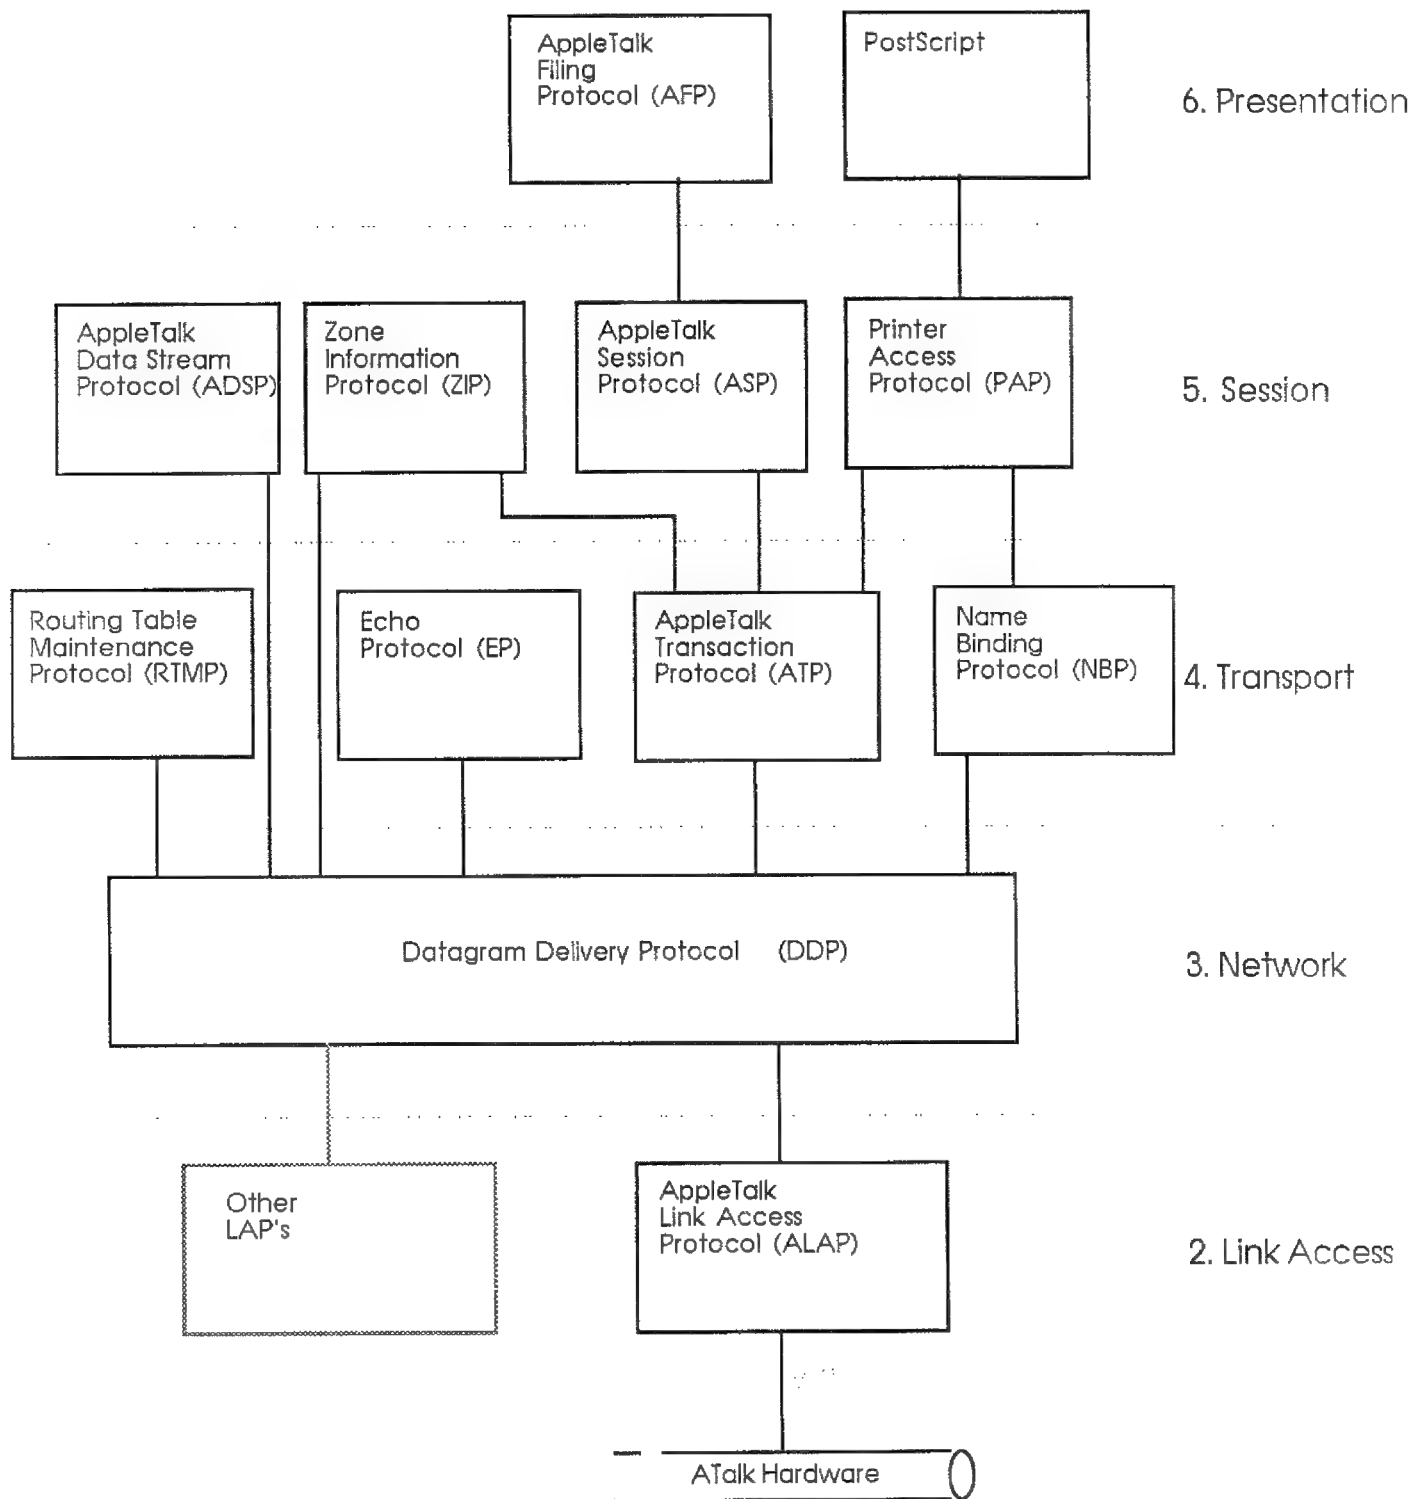


Figure I-1. AppleTalk Protocols: A Summary

II. Electrical Specifications

The electrical and mechanical characteristics of AppleTalk correspond to the Physical layer of the ISO-OSI model. AppleTalk is a multi-drop, balanced, transformer-isolated serial communications system. The raw data rate is 230.4 kilobits per second over a maximum distance of 300 meters.

SDLC frame format is used, and FM-0 modulation (FM-0 is a bit-encoding technique that provides self-clocking). Balanced signalling is achieved using RS-422 drivers and receivers in each of the attached devices. The transformer provides ground isolation as well as protection from static discharge. Since devices are passively connected to the trunk cable via a drop cable, a device may fail without disturbing communications. Devices can be added and removed from the bus with only minor disruption of service.

The Physical layer performs the functions of bit encoding/decoding, signal transmission/reception, and carrier sense; these functions are discussed below.

Bit Encoding and Decoding

Bits are encoded using a self-clocking technique known as FM-0 (bi-phase space). In FM-0, each bit cell (nominally, 4.34 microseconds) contains a transition at its end, thus providing timing information (known as one *bit-time*). Zeros (0's) are encoded by adding an additional transition at mid-cell (see figure II-1).

Signal Transmission and Reception

The use of the EIA RS-422 signalling standard for transmission and reception over AppleTalk provides significantly higher data rates over longer distances than that of the EIA RS-232C standard. AppleTalk uses differential, balanced voltage signalling at 230.4 kilobits per second over a maximum distance of 300 meters. The balanced configuration provides better isolation from ground noise currents, and is not susceptible to fluctuating voltage potentials between system grounds or common-mode electromagnetic interference (EMI).

Carrier Sense

The Physical layer provides an indication to ALAP when activity is sensed on the cable. Two indications are provided: SDLC frame in progress (the *hunt bit*), and missing clock detected. The hunt bit indicates when the hardware is searching for the start of the next SDLC frame - when this bit is cleared, the hardware is in the middle of an SDLC frame. Note that a frame can not be detected until a complete flag has been sent on the line and recognized by the hardware. Missing clock, on the other hand, works in conjunction with the synchronization pulse sent before certain frames (see chapter III), to provide a more immediate indication of an ongoing transmission. Technically, missing clock indicates the detection, and then the absence, of a clocking signal on the line.




Figure II-1. FM-0 Encoding

III. AppleTalk Link Access Protocol (ALAP)

About AppleTalk Link Access Protocol

The AppleTalk Link Access Protocol (ALAP), which corresponds to the data link layer of the ISO-OSI reference model, is the key mechanism allowing AppleTalk devices to share the communication medium (in this chapter also referred to as the *bus*). ALAP, which must be common to all systems/devices on the bus, provides best effort delivery of information between these systems/devices, also known as *nodes*. It provides the basic underlying service of packet transmission between the nodes connected to a single AppleTalk network. ALAP manages the encapsulation of data in a *frame* and then provides access to the bus for transmission and reception of frames. The main goals of the data link protocol for a shared-bus system like AppleTalk are as follows:

- provide bus access management
- provide a node addressing mechanism
- perform data transmission and reception
- ensure packet length and integrity

Bus Access Management

All AppleTalk nodes compete for use of the bus. It is the function of ALAP to resolve this contention and provide fair access for all nodes. ALAP uses an access protocol known as *Carrier Sense Multiple Access with Collision Avoidance* (CSMA/CA). *Carrier sense* means that a node sending data first senses the line, and defers to any ongoing transmission. *Collision avoidance* means that the protocol attempts to minimize collisions. A collision occurs when two or more nodes transmit frames at the same time. In the AppleTalk CSMA/CA technique, all transmitters wait until the line has been idle for a minimum time, plus an additional time as determined by the generation of pseudo-random numbers whose range is adjusted according to the perceived bus traffic. AppleTalk hardware does not allow the actual detection of collisions.

Node Addressing Mechanism

ALAP uses an 8-bit *node identifier number* (also known as the *node ID*) for identifying nodes on the bus. ALAP is ultimately responsible for encoding both the destination and source node IDs in the header portion of an outgoing frame. The destination node ID is used to *filter* frames at the data link layer; frames whose destination node ID does not match the receiving node's ID, or the broadcast ID (\$FF) are not accepted at that node.

Unlike other networks that use fixed, universally-unique addresses, AppleTalk uses a dynamic node address assignment scheme.

Dynamic Node ID Assignment

A key motivation for dynamic node ID assignment is that when a node is moved between networks, its old node ID should not conflict with a node ID already in use on the new network. Furthermore, the availability of such a scheme eliminates one part of network configuration. Unlike networks such as Ethernet no special universally-unique number

need be built into each node, nor is it necessary to administer the assignment of such numbers to different vendors.

When a node is activated on an AppleTalk network, it makes a guess at its own node ID, either by extracting this number from some form of long-term (for example, parameter RAM or disk) memory, or by generating a random number (if the node has no nonvolatile memory). The new node must then verify that this guessed number is not already in use on that network.

It does this by sending out a special frame, known as an ALAP *enquiry control frame*, to the guessed node address and waits for acknowledgement. If the guessed node ID is in use, then the corresponding node that is using it will, upon receiving the enquiry, respond with an ALAP *acknowledge control frame*. The reception of this frame tells the new node that the guessed node ID is already in use, and that the process must be repeated with a different guess. Each enquiry control frame is transmitted repeatedly to account for cases where a node already using the guessed node ID might be busy, and thus miss an enquiry frame.

Node ID numbers are divided into two classes: *server node IDs* and *user node IDs*. User node IDs are in the range 1 to 127 and server node IDs are in the range 128 to 254. A destination node ID of 255 has a special meaning. Frames received with the destination node ID equal to this value are accepted by all nodes; this permits the *broadcasting* of packets to all nodes on the network. A destination node ID equal to 0 is not allowed, and is treated as *unknown*.

This division of node IDs is due to the fact that some nodes may, for extended periods of time, disable reception from AppleTalk (for instance, if they are engaged in a device-intensive operation such as disk access or transferring a bit map document to a laser printer). Such a node would not respond to another node's enquiry frames, which could result in two nodes acquiring the same node ID.

It is extremely important that no node acquire the same address as an already-functioning server node; this would disrupt service not only for the conflicting nodes, but for other users of the server as well. Excluding user (that is, non-server) node IDs from the server node ID range eliminates the possibility of user nodes (which are switched on and off with greater frequency) conflicting with servers.

Within the user node ID range, verification can be done more expeditiously (that is, fewer retransmissions of the enquiry frame) to decrease the initialization time for such nodes. A more thorough node ID assignment scheme is used by servers (that is, additional time taken, once chosen, to ensure that they will be unique on the network). This is not detrimental to the server's operation because such nodes are rarely switched on and off.

If during normal operation, after acquiring its node ID, a node detects the use of the same ID by some other node (by receiving a Clear-to-Send control frame with a source node ID identical to its own ID), ALAP should set a flag informing its client of this conflict condition. Recovery from this condition, if any, is the responsibility of higher-level protocols.

Data Transmission and Reception

AppleTalk uses a bit-oriented data link protocol for data transmission and reception. Unlike byte-oriented protocols, a bit-oriented protocol permits the use of all possible bit

patterns within the frame. The frame delimiter for ALAP, known as a *flag* byte, is the distinguished bit sequence 01111110 (\$7E). Typically, flags are generated by (hardware) transmitters at the beginning and end of frames; they are used by (hardware) receivers to detect frame boundaries.

In order for the data link protocol to transmit all possible bit patterns within a frame, the protocol must insure *data transparency*. ALAP does this with a technique known as "bit stuffing". When transmitting a frame, after each string of five consecutive 1's detected in the user data stream, it inserts a 0; this guarantees that data sent on the bus contains no sequences of more than five consecutive 1's. A receiving ALAP performs the inverse operation, "stripping" a 0 that follows five consecutive 1's.

Prior to transmitting a frame, ALAP sends out a synchronization pulse followed by a *frame preamble* consisting of two or more flags bytes. The frame is followed at its end by a *frame postamble/trailer* consisting of a flag byte and an abort sequence (seven or more 1 bits).

ALAP Frame Format

An ALAP frame consists of a three-byte ALAP header followed by a variable length data field (0 to 600 bytes), and a 16-bit frame check sequence. Its format is summarized in figure III-1.

ALAP Header

The ALAP header contains the destination node ID, the source node ID, and a one-byte *ALAP type field*. The ALAP type field specifies the type of frame. Values in the range 128 to 255 (\$80 to \$FF) are reserved for *ALAP control frames*; these frames do not contain a data field. The currently used ALAP control frames are as follows:

- *lapENQ* (ALAP type \$81): an enquiry control frame, used in dynamic node ID assignment.
- *lapACK* (ALAP type \$82): an acknowledgement control frame, sent in response to a *lapENQ*.
- *lapRTS* (ALAP type \$84): a Request-to-Send control frame that notifies the destination node that the transmitter wants to send a data packet to it.
- *lapCTS* (ALAP type \$85): a Clear-to-Send control frame sent in response to a *lapRTS*, indicating the willingness of the receiver to receive a data packet.

ALAP control frames received with values in the ALAP type field other than those listed above are simply discarded.

ALAP type fields with values in the range 1 to 127 (\$01 to \$7F) are used for *ALAP data frames*; these frames carry client data in the data field. In such frames, the type field is used to specify the ALAP protocol type of the client to whom the frames's data must be delivered. This allows the concurrent use of ALAP by several network layer protocols and is crucial to maintaining an "open systems architecture". The ALAP implementation in the receiving node uses the type field to determine the client for whom the data is intended. This client in turn uses this field to decide how to interpret the data (format of the data,

header for a higher-level protocol, etc.). As an example, the Datagram Delivery Protocol uses the values 1 and 2 in the ALAP type field.

ALAP simply transmits and receives data packets on behalf of its clients. The format and interpretation of the data field is defined by higher-level protocols, with the following exception: the low-order ten bits of the first two bytes of the data field must contain the length in bytes of the data field, more significant bits first (the data length includes the length field itself).

Frame Check Sequence Field

The 16-bit frame check sequence (FCS) is computed as a function of the contents of the destination node ID, source node ID, LAP type, and data fields, using the standard CRC-CCITT cyclic redundancy check algorithm. This is described in detail in the Algorithm section of this chapter.

Frame Size Limitations

Since the ALAP header is 3 bytes and the data field can contain from 0 to 600 bytes, the smallest valid frame (not including the preamble, postamble, and the FCS) is 3 bytes long, while the largest is 603 bytes.

Frame Transmission

The transmission of a data frame by ALAP involves a special *dialogue* consisting of one or more ALAP control frames followed by the data frame. This dialogue is based on a CSMA/CA access protocol some aspects of which were outlined above in the section "Bus Access Management." The exact form of the dialogue depends on the destination of the frame; on this basis, ALAP distinguishes two kinds of frames: *directed* and *broadcast*. A directed packet is one whose destination address is a single node; a broadcast packet is intended to be received by all nodes.

Dialogues must be separated by a minimum *Inter-Dialogue Gap* (IDG) of 400 microseconds; the different frames of a single dialogue must follow one another with a maximum *Inter-Frame Gap* (IFG) of 200 microseconds.

The frame transmission procedure is described separately for directed and for broadcast frames. Figure III-2 diagrams the packets and timing used in this procedure.

Directed Data Frames

The transmitting node uses the physical layer's ability to sense if the line is in use. If the line is busy the node waits until it becomes idle following the current ongoing transmission. The node is said to *defer*. Upon sensing an idle line, the transmitter waits for a time equal to the minimum IDG plus a randomly generated amount. During this "wait," the transmitter continues to monitor the line. If the line becomes busy at any time during this wait period, the node must again defer. If the line remains idle throughout this wait period, then the node sends a *lapRTS* frame to the intended receiver of the data frame. The receiver must, within the maximum IFG, return a *lapCTS* frame to the transmitting node. Upon receiving this frame, the transmitter must within the maximum IFG send out the data frame.

The purpose of this algorithm is twofold: (1) to restrict the periods in which collisions are highly likely (this is during the *lapRTS-lapCTS* exchange), and (2) to spread out in time several transmitters waiting for the line to become idle. The *lapRTS-lapCTS* exchange, if successfully completed, signifies that a collision did not occur, and that all intending transmitters have heard of the coming data frame transmission and are deferring/waiting.

If in fact a collision does occur during the *lapRTS-lapCTS* exchange, a *lapCTS* will not be received, and the sending node will then *back off* and retry. The sending node is said to presume a collision.

The range of the random wait time generated by a node is adjusted if it presumes such a collision. The idea is that if collisions have been presumed for recently sent packets, this signifies heavier bus loading and higher contention for the bus. In this case, the random wait should be generated over a larger range, thus spreading out (in time) the different contenders for the line. Conversely, if the node has not had to defer on recent transmissions, a lighter load is signified and the random wait should be generated over a smaller range, thus reducing transmission time.

Two factors are used for adjusting the range: (a) the number of times the node had to defer, and (b) the number of times it presumed a collision. This history is maintained in two 8-bit *history bytes*, one each for deferences and collisions. At each attempt to send a packet these bytes are shifted left one bit. The lowest bit of each byte is then set if the node had to defer or presume a collision, respectively, on that attempt, else this bit is cleared. In effect, the history bytes remember the deference and collision history for the last eight attempts. The exact use of these bytes for determining the random wait times is described in the Algorithm section.

Synchronization

At the beginning of all RTS frames, ALAP transmits a *synchronization pulse*. This is a transition on the bus, followed by an idle period greater than two bit-times. The synchronization pulse is obtained by momentarily enabling the line driver for at least one bit-time, before disabling it. This causes a transition that will be taken as a clock by all receivers on the network. However, since it is followed by an idle period of sufficient length, a *missing clock* is detected by the receivers. The missing clock allows transmitters to synchronize their access to the line (they become immediately aware if someone is about to transmit). Sync pulses may also be optionally sent at the beginning of other LAP frames.

Directed Transmissions

Directed Packets are sent via a 3-frame dialogue as follows:

- (1) The transmitter senses the bus until the bus is idle for the minimum IDG time.
- (2) The transmitter waits an additional time determined by a random number.
- (3) The transmitter sends a *lapRTS* frame to the intended destination node.
- (4) The receiver sends a *lapCTS* frame to the transmitter.
- (5) The transmitter, upon successful reception of the *lapCTS*, sends a data frame (in which it encapsulates the client's packet).

Note: all response frames must be sent within the maximum IFG time.

The extra wait during step 2 tends to spread out transmitters which are contending for the line and, thus, minimize (or avoid) collisions. A transmitter presumes that a collision has occurred when it does not receive a CTS frame in the required time (IFG). When this happens, the transmitter retries starting at step 1. For each attempt, a new random number must be generated in step 2. If after 32 attempts the transmitter is unable to send the data frame, it reports failure to its client.

Broadcast Transmissions

Broadcast packets are distinguished by their destination node ID being 255 (\$FF). When ALAP wishes to send a broadcast packet, it performs the following procedure:

- (1) The transmitter senses the bus until the bus is idle for the minimum IDG time.
- (2) The transmitter waits an additional time determined by a random number.
- (3) The transmitter sends a *lapRTS* frame with destination address \$FF.
- (4) The transmitter senses the line for the maximum IFG time.
- (5) After sensing the line idle for IFG, it sends its DATA frame.

The purpose of sending the *lapRTS* in step 3 is to ensure that other transmitters are cognizant of the intent to transmit, and to force a collision if another transmitter starts at the same time (that transmitter will then backoff). If the transmitter detects bus activity during step 4, it makes up to 32 further attempts beginning with step 1. If it fails to transmit the data frame after 32 attempts, it reports failure to its client.

Broadcast packets are sent without collision except in the unlikely case of another transmitter attempting a broadcast at the same time.

Frame Reception

Frame reception is in many ways the inverse of the frame transmission process: frames are decapsulated and accepted only if their destination address is equal to the node's (or is equal to the broadcast address of 255) and their CRC is verified. The received frame can generate several conditions that must be handled by ALAP: overrun error, bad frame size, underrun error, bad frame type, and bad frame CRC.

Frame Size Error

If a frame (excluding the FCS), of greater than 603 bytes or smaller than 3 bytes is received then it is rejected and a bad frame error status generated.

Overrun/Underrun Error

If ALAP was not able to stay synchronized with the incoming data, causing receiver overruns or underruns, the frame will be rejected and an overrun or underrun error generated.

Frame Type

The frame's ALAP type field is checked against each of the valid values for an ALAP frame. If none match, the frame is rejected, and a bad frame-type error is generated.

Frame Check Sequence

The CRC-CCITT frame check sequence is computed for the incoming frame. When the entire frame has been received, ALAP tests the computed FCS for the frame. If the CRC is in error the frame is rejected and a bad frame-CRC error is generated.

Algorithms

CRC-CCITT Calculation

The CRC-CCITT calculation algorithm uses the standard CRC-CCITT polynomial:

$$G(x) = x^{16} + x^{12} + x^5 + 1$$

The CRC-CCITT frame check sequence value corresponding to a given frame is calculated based on the following polynomial division identity:

$$\frac{M(x)}{G(x)} = Q(x) + \frac{R(x)}{G(x)}$$

where:

$M(x)$ = binary polynomial (corresponding to the frame after complementing its first 16 bits)

$R(x)$ = remainder after dividing $M(x)$ by the generating polynomial (its coefficients are the bits of the CRC)

$Q(x)$ = quotient after dividing $M(x)$ by the generating polynomial (its coefficients are ignored when calculating the CRC)

The implementation of the CRC for the FCS field, at the transmitter, computes the CRC starting with the first bit of the destination node ID following the opening flag and stopping at the end of the data field. The FCS field is equal to the ones-complement of the transmitter's remainder. The same calculation is performed at the receiver, however the FCS bytes themselves are included in the computation. The result of a correctly received transmission is then the binary constant 0001110100001111 ($x^{15} \dots x^0$). This constant is a characteristic of the divisor.

In addition to the division of the binary value of the data by the generating polynomial to generate the remainder for checking, the following manipulations occur:

- The dividend is initially preset to all ones before the computation begins; this has the effect of complementing the first 16 bits of the data.
- The transmitter's remainder is inverted bit-by-bit (FCS field) as it is sent to the receiver. The high-order bit of the FCS field is transmitted first ($x^{15} \dots x^0$).

If the receiver computation does not yield the binary constant 0001110100001111, the frame is discarded.

Random Wait Time Determination

The range of the random wait before attempting to send a frame after the line becomes idle is a function of the past history of attempted transmissions. As described in the "Frame Transmission" section, two history bytes are maintained indicating the number of presumed collisions and the number of deferrals in the last eight transmission attempts. These two bytes are used to adjust a *global backoff mask*. The mask takes on particular values between \$00 and \$0F (specifically binary 0, 1, 11, 111 and 1111). These values determine the range of random number picked, as specified below. The global backoff mask is adjusted at the beginning of a request to transmit a client's data as follows:

- The mask starts as 0
- If the number of times the node backed off during the last 8 attempts is greater than 2, the mask is extended by one bit up to the maximum of 4 bits (logical shift left and set the low-order bit) and the backoff history byte is then cleared.
- Else, if the number of times the node had to defer is less than 2, the mask is reduced by one bit (logical shift right) and the defer history byte is set to all one's.
- Else, if neither of these apply, the mask is left as is.

Due to collisions and deferrals, ALAP may have to make many attempts to send a packet. The following operations are performed, in order, before making the first attempt:

- The global backoff mask is adjusted as specified above.
- The two history bytes are shifted left one bit and the low-order bit of each is set to zero.
- The global backoff mask is copied into a *local backoff mask*.

During each attempt to send a packet, the following operations are performed, in order:

- If the line is busy, the node defers until the end of the dialogue. The low-order bit of the deferral history byte is set to one, and the low-order bit of the local backoff mask is also set to one (in case the mask was zero).
- If the line is not busy, the node waits 400 microseconds. If the line becomes busy during this time, the node defers as above.
- The random wait time is generated: a random number is picked and masked by (ANDed with) the local backoff mask. This random number thus has a range of zero to fifteen. The node waits for 100 microseconds multiplied by this random number. If the line becomes busy during this time, the node defers.
- The node sends a *lapRTS*. If a *lapCTS* is received within the maximum IFG (or if the packet is to be broadcast), the data is sent. If not (a collision is presumed), the low-order bit of the collision history byte is set to one, and the local backoff mask is shifted left one bit and its low-order bit set. The node then tries again.
- If, during an attempt to send a packet, 32 collisions occur or the node has to defer 32 times, the attempt is aborted and an error returned to the ALAP client.

ALAP Procedural Model

The following procedural model, written in a Pascal-like language, is provided as a specification of ALAP. Any particular ALAP implementation's logical and timing behavior should follow this model.

Assumptions

The model assumes that the program executes sufficiently fast so as to not introduce any execution delay into the timing of events. Where the ALAP specifies a timing delay, this is assumed to be performed via references to a global real function RealTime which returns the current time in microseconds. All timing constraints are specified as real constants.

The model assumes sequential, single-process execution of the code. This is especially true of TransmitPacket and ReceivePacket (and the procedures they call).

In a typical implementation, packet reception will be triggered by means of a hardware interrupt. The interrupt routine will then execute the ReceiveLinkMgmt procedure. The interrupt routine must provide a mechanism for saving valid data packets and for informing higher level protocols of this event. Such specifications are outside the scope of this document.

It is also assumed that a transmitter continuously listens to the bus while it is waiting for its access to the line.

Global Constants, Types and Variables

The following global constants, types and variables are used throughout the model.

```
CONST
    minFrameSize = 3;           { smallest (LAP header only) frame }
    maxFrameSize = 605;        { size of largest LAP frame including FCS }
    maxDataSize = 600;         { size of largest (encapsulated) ALAP data field }
    bitTime = 4.34;            { bit time (μsec) }
    byteTime = 39.0;           { worst case single byte time (μsec) }
    minIDGtime = 400.0;         { minimum Inter-Dialog Gap (μsec) }
    IDGslottime = 100.0;        { slot time of transmit backoff algorithm (μsec) }
    maxIFGtime = 200.0;         { maximum Inter-Frame Gap (μsec) }
    maxDefers = 32;             { maximum defers for a single packet }
    maxCollsns = 32;            { maximum collisions for a single packet }
    lapENQ = $81;               { LAP type field value of ENQuiry frame }
    lapACK = $82;               { ... ACKnowledgement frame }
    lapRTS = $84;               { ... ReqeustToSend frame }
    lapCTS = $85;               { ... ClearToSend frame }
    hdlcFLAG = $7E;             { value of an HDLC FLAG }
    wksTries = 20;              { Number of ENQ sets for a workstation to try }
```

TYPE

```
{ global result types from LAP functions }
TransmitStatus = (transmitOK, excessDefers, excessCollsns, dupAddress);
```

```
ReceiveStatus = (receiveOK, Receiving, nullReceive, frameError);
FrameStatus = (noFrame, lapDATAframe, lapENQframe, lapACKframe, lapRTSframe,
               lapCTSframe, badframeCRC, badframeSize, badframeType,
               overrunError, underrunError);

{ Data Link types and structures }
bit = 0 .. 1;
bitVector = PACKED ARRAY [0 .. 7] OF bit;
octet = $00 .. $FF;
anAddress = octet;
aLAPtype = octet;
aDataField = PACKED ARRAY [1 .. maxDataSzie] OF octet;

{ Basic structure of an ALAP frame, not including FLAGS, FCS }
frameInterpretation = (raw, structured);
aFrame = PACKED RECORD
    CASE frameInterpretation OF
        raw: rawData : PACKED ARRAY [1 .. maxFrameSize] OF octet;
        structured: (
            dstAddr : anAddress;
            srcAddr : anAddress;
            lapType : aLAPtype;
            dataField : a DataField)
    END;

VAR
    MyAddress : octet;           { set during initializeLAP }
    Backoff : INTEGER;           { current backoff range }
    fAdrValid,                   { MyAddress has been validated }
    fAdrInUse,                   { Another node has same MyAddress }
    fCTSexpected : BOOLEAN;      { RTS has been sent, CTS is valid }
    deferCount, collsnCount : INTEGER; { optional, for statistics only }
    deferHistory, collsnHistory : bitVector;
    outgoingLength, incomingLength : INTEGER;
    outgoingPacket, incomingPacket : aFrame;
```

Hardware Interface Declarations

The following declarations refer to hardware-specific interfaces which are assumed to be available to the LAP procedures. The functions are typically bits and/or bytes contained in the relevant hardware interface chip(s). Similarly, the procedures are expected to be represented in actual hardware by means of control bits within the hardware.

A brief description of the assumed attributes of each of these follows.

CarrierSense indicates that the hardware is sensing a frame on the bus
RcvDataAvail indicates that a data byte is available
rxDATA is the next data byte available (as indicated by RcvDataAvail)
EndOfFrame indicates that a valid closing FLAG has been detected
CRCok indicates that the received frame's FCS is correct (when EndOfFrame)
OverRun indicates that the code did not keep up with data reception
MissingClock indicates that a missing clock has been detected
setAddress sets the hardware to receive frames directed to MyAddress
enableTxDrivers and disableTxDrivers control the operation of the RS-422 drivers
enableTx and disableTx control the operation of the data transmitter
txFLAG causes the transmission of a FLAG

txDATA causes the transmission of a data byte
txFCS causes the transmission of the Frame Check Sequence
txONES causes 12+ one (1's) to be transmitted
resetRx, enableRx and disableRx control the receiver
resetMissingClock causes the MissingClock indication to be cleared

```
{ Hardware interface functions/procedures }
FUNCTION CarrierSense : BOOLEAN;    EXTERNAL;
FUNCTION RxDataAvail : BOOLEAN;    EXTERNAL;
FUNCTION rxDATA : octet;            EXTERNAL;
FUNCTION EndOfFrame : BOOLEAN;    EXTERNAL;
FUNCTION CRCok : BOOLEAN;          EXTERNAL;
FUNCTION OverRun : BOOLEAN;        EXTERNAL;
FUNCTION MissingClock : BOOLEAN;    EXTERNAL;
PROCEDURE setAddress (addr : octet); EXTERNAL;
PROCEDURE enableTxDrivers;          EXTERNAL;
PROCEDURE disableTxDrivers;        EXTERNAL;
PROCEDURE enableTx;                EXTERNAL;
PROCEDURE txFLAG;                  EXTERNAL;
PROCEDURE txDATA (data : octet);    EXTERNAL;
PROCEDURE txFCS;                   EXTERNAL;
PROCEDURE txONES;                  EXTERNAL;
PROCEDURE disableTx;               EXTERNAL;
PROCEDURE resetRx;                 EXTERNAL;
PROCEDURE enableRx;                EXTERNAL;
PROCEDURE disableRx;               EXTERNAL;
PROCEDURE resetMissingClock;        EXTERNAL;
```

Interface Procedures and Functions

The ALAP model's interface to the next higher layer (its client) is specified in terms of the following three calls:

(1) PROCEDURE initializeLAP (hint : octet; server : BOOLEAN)

This procedure initializes the ALAP; it is expected to be called exactly one. The hint parameter is a suggested starting value for the node's AppleTalk physical link address; a value of zero indicates that ALAP should generate a starting value. Upon return from the call, the station's actual address is available in the global variable **MyAddress**. If **server** is true then the internal procedure **acquireAddress** will spend extra time to determine if another node is using the selected node address.

(2) FUNCTION transmitPacket (dstParam : anAddress; typeParam : aLAPtype;
dataField : aDataField; dataLength : INTEGER) : TransmitStatus;

This is the call provided to transmit a packet. The internal function **TransmitLinkMgmt** performs the transmission Link Access algorithms.

(3) PROCEDURE receivePacket (VAR dstParam : anAddress;
VAR srcParam : anAddress; VAR typeParam : aLAPtype;
VAR dataField : aDataField; VAR dataLength : INTEGER);

This is the entry provided to receive a packet. The internal function `ReceiveLinkMgmt` implements the reception Link Access algorithms.

InitializeLAP Procedure

`initializeLAP` is called to reset the global variables to known states; it calls `acquireAddress` to initialize `MyAddress`.

```
PROCEDURE initializeLAP (hint : octet; server : BOOLEAN);  
  
VAR i : INTEGER;  
  
BEGIN  
    Backoff := 0;  
  
    { initialize history data for Backoff calculations }  
    FOR i := 0 TO 7 DO  
        BEGIN  
            deferHistory [i] := 0;  
            collsnHistory[i] := 0;  
        END;  
  
    deferCount := 0; collsnCount := 0;    { optional }  
  
    acquireAddress (hint, server);  
END;
```

AcquireAddress Procedure

The procedure `acquireAddress` provides the dynamic node assignment algorithm. A special frame (of type `lapENQ`) is created and sent. When no node responds after repeated attempts, the current value of `MyAddress` is assumed to be safe for use by this node; the state of `fAdrValid` reflects this fact. If the global `fAdrInUse` ever becomes true after a call to `AcquireAddress`, another node that is using the same `MyAddress` has been detected.

```
PROCEDURE AcquireAddress (hint : octet; server : BOOLEAN);  
  
VAR maxTrys, try : INTEGER; ENQframe : aFrame;  
  
BEGIN  
    IF hint > 0 THEN myAddress := hint  
    ELSE IF server THEN MyAddress := Random (127) + 128  
    ELSE MyAddress := Random (127) + 1;  
  
    setAddress (MyAddress);  
    fAdrValid := FALSE;  
    IF server THEN maxTrys := wksTries * 6 { servers try six times as much as workstations }  
    ELSE maxTrys := wksTries;  
  
    try := 0; fAdrInUse := false;  
  
    { the main loop of acquireAddress -- repeatedly check for response to ENQ }  
  
    WHILE try < maxTrys DO
```

```
IF (transmitPacket (MyAddress, lapENQ, ENQframe.dataField,0) = transmitOK) OR
  fAdrInUse THEN BEGIN
  IF server THEN MyAddress := Random (127) + 128
  ELSE MyAddress := Random (127) + 1;
  setAddress (MyAddress);
  trys := 0;
  END { IF }
  ELSE trys := trys + 1;
END; { WHILE }

fAdrValid := TRUE;
END; { AcquireAddress }
```

TransmitPacket Function

The function `transmitPacket` is called by the ALAP client to send a packet. After constructing (encapsulating) the caller's `dataParam`, it calls upon `TransmitLinkMgmt` to perform the actual link access.

```
FUNCTION transmitPacket (dstParam : anAddress; typeParam : aLAPtype;
  dataParam : aDataField; dataLength : INTEGER) : TransmitStatus;

BEGIN
  IF fAdrInUse THEN transmitPacket := dupAddress
  ELSE BEGIN

    { copy interface data into frame for TransmitLinkMgmt }
    WITH outgoingPkt DO BEGIN
      dstAddr := dstParam;
      srcAddr := MyAddress;
      lapType := typeParam;
      dataField := dataParam;
    END; { WITH }

    outgoingLength := dataLength + 3;
    transmitPacket := TransmitLinkMgmt;
  END; { ELSE }
END; { transmitPacket }
```

TransmitLinkMgmt Function

The function `TransmitLinkMgmt` implements the Carrier Sense, Multiple Access with Collision Avoidance algorithm; `TransmitLinkMgmt` is at the heart of the AppleTalk Link Access Protocol.

The typical AppleTalk hardware is not capable of performing collisions detection.

ALAP attempts to minimize collisions by requiring transmitters to wait a randomly generated amount of time before sending their RTS frames after the bus has been sensed idle for a minimum time (the IDG). Any transmitter which detects that another transmitter is in progress while it is in this random wait must defer.

In order to minimize delays under light loading, but still be able to minimize the probability of collisions under moderate to heavy loading, the random delay is picked in a range that is

constantly adjusted based upon the recently observed history of a node's attempts to access the bus. Two history vectors (`deferHistory` and `collsnHistory`) are used to keep track of the last 8 access attempts at the bus. Each attempt adds a single bit of history data to these vectors by shifting left the current values and updating the vacated bit with the appropriate value. `DeferHistory[0]` is set if a deferral is required during an access; `collsnHistory[0]` is set if a collision is assumed.

The range of the current Backoff is adjusted upwards (to a maximum of 16) when the observed collisions exceeds 2 in the last 8 attempts; Backoff is adjusted downwards if the number of observed deferrals is less than 2 in the last 8 attempts. When an adjustment is made, the corresponding history data is set to the "maximum" value so that further adjustments are inhibited until more history data has accumulated. The maximum value for `deferHistory` is all 1's; the maximum value for `collsnHistory` is all 0's.

Note that a local backoff (`LclBackoff`) is used during the retry attempts of a given frame. This has the effect of spreading out attempts to a non-listening node for a longer time, thus increasing its chances of receiving the packet.

```
FUNCTION TransmitLinkMgmt : TransmitStatus;
```

```
VAR
```

```
  LclBackOff, i : INTEGER;  
  fBroadcast, fENQ : BOOLEAN;  
  xmttimer : REAL;  
  rcvdfame : FrameStatus;  
  RTSframe : aFrame;
```

```
BEGIN
```

```
  WITH RTSframe DO BEGIN
```

```
    dstAddr := outgoingPacket.dstAddr;  
    srcAddr := MyAddress;  
    lapType := lapRTS
```

```
  END;
```

```
  fBroadcast := (outgoingPacket.dstAddr = $FF);
```

```
  fENQ := (outgoingPacket.lapType = lapENQ);
```

```
  { Adjust Backoff, based upon recent history }
```

```
  IF bitCount (collsnHistory) > 2 THEN BEGIN
```

```
    Backoff := min (max (Backoff * 2, 2), 16);
```

```
    FOR i := 0 TO 7 DO collsnHistory [i] := 0;
```

```
  END      { IF }
```

```
  ELSE IF bitCount (deferHistory) < 2 THEN BEGIN
```

```
    Backoff := Backoff DIV 2;
```

```
    FOR i := 0 TO 7 DO deferHistory[i] := 1
```

```
  END;      { ELSE IF }
```

```
  { Shift history data }
```

```
  FOR i := 7 DOWNT0 1 DO BEGIN
```

```
    collsnHistory [i] := collsnHistory [i-1];
```

```
    deferHistory [i] := deferHistory [i-1];
```

```
  END;      { FOR }
```

```
  collsnHistory [0] := 0; deferHistory [0] := 0;
```

```

{ Initialize main loop }
deferTries := 0; collsnTries := 0;
LclBackoff := Backoff; transmitdone := FALSE;

{ Begin main loop }
REPEAT

    { Wait for minimum Inter-Dialog Gap time }
    REPEAT

        { Wait for any packet in progress to pass }
        IF CarrierSense THEN BEGIN
            { Ensure minimum backoff if packet in progress }
            LclBackoff := max (LclBackoff, 2);
            deferHistory [0] := 1;

            { Perform watchdog reset of Rx for "stuck" CarrierSense }
            xmttimer := RealTime + 1.5 * maxFrameSize * byteTime;
            REPEAT UNTIL ( NOT CarrierSense) OR (RealTime > xmttimer);
            IF CarrierSense THEN ResetRx;          { something's wrong, clear it }
        END; { IF }

        { Wait for minimum IDG after packet (or idle line) }
        xmttimer := RealTime + minIDGtime;
        REPEAT UNTIL (RealTime > xmttimer) OR CarrierSense;

    UNTIL NOT CarrierSense;

    ResetMissingClock;

    { Wait our additional backoff time, deferring to others }
    xmttimer := RealTime + Random (LclBackoff) * IDGslottime;
    REPEAT UNTIL (RealTime > xmttimer) OR CarrierSense;

    IF CarrierSense OR MissingClock THEN BEGIN { defer }
        DeferCount := DeferCount + 1;          { optional }
        LclBackoff := max (LclBackoff, 2);
        deferHistory [0] := 1;
        IF deferTries < maxDefers THEN deferTries := deferTries + 1
        ELSE BEGIN
            TransmitLinkMgmt := excessDefers;
            transmitdone := TRUE;
        END { ELSE }
    END { IF }

    ELSE BEGIN { NOT (CarrierSense OR MissingClock) }
        IF fENQ THEN transmitFrame (outgoingPacket, 3)
        ELSE transmitFrame (RTSFrame, 3);

        { use common code to detect line state }
        fCTSexpected := TRUE;
        rcvdfame := receiveFrame;
        fCTSexpected := FALSE;

        IF fAdrInUse THEN BEGIN
            TransmitLinkMgmt := dupAddress;
            transmitDone := TRUE;
        END { IF }
    END

```

```
ELSE CASE rcvdFrame OF
  noFrame: IF fBroadcast THEN BEGIN
    transmitFrame (outgoingPacket, outgoingLength);
    transmitLinkMgmt := transmitOK;
    transmitdone := TRUE;
  END;

  lapCTSframe : IF (NOT fENQ) AND (NOT fBroadCast) THEN BEGIN
    transmitFrame (outgoingPacket, outgoingLength);
    transmitLinkMgmt := transmitOK;
    transmitdone := TRUE;
  END;
END;      { CASE }

{ Assume collision if we don't receive the expected CTS }
IF NOT transmitdone THEN BEGIN
  CollsnCount := CollsnCount + 1;      { optional }
  collsnHistory[0] := 1;                { update history data }
  IF collsnTries < maxCollsns THEN BEGIN
    LclBackoff := min (max (LclBackoff*2,2),16);
    collsnTries := collsnTries + 1;
  END      { IF }
  ELSE BEGIN
    TransmitLinkMgmt := excessCollsns;
    transmitdone := true;
  END      { ELSE }
END      { IF NOT ... }
END      { ELSE NOT ... }

UNTIL transmitdone;
END;      { TransmitLinkMgmt }
```

TransmitFrame Procedure

The procedure transmitFrame is responsible for putting data on the bus. Notice that certain details, such as how a FLAG is forced and a packet terminated (which includes sending of the FCS) are not explicitly detailed here due to their extreme hardware dependence. Note: the 12 ones at the end of the frame are required to finish clocking of the data by a receiver.

```
PROCEDURE transmitFrame (VAR frame : aFrame; framesize : INTEGER);

VAR
  i : INTEGER;
  bittimer : REAL;

BEGIN

  disableRx;

  { Generate the synchronizing pulse -- really only required before RTS frames }
  bittimer := RealTimer + 1.5 * bitTime;
  enableTxDrivers;
  WHILE RealTime < bittimer DO BEGIN END;
  disableTxDrivers;
  bittimer := RealTimer + 1.5 * bitTimer;
```

```

WHILE RealTime < bittimer DO BEGIN END;

{ Start the actual frame transmission }
enableTxDrivers;
enableTx;
txFLAG; txFLAG;          { Output 2 opening FLAG's }
FOR i := 1 to framesize DO TxData (frame.rawData[i]);
txFCS;                    { Send the FCS }
txFLAG;                   { Send the trailing FLAG }
txONES;                   { Send 12 1's for extra clocks }
disableTxDrivers;

{ re-establish default listening mode }
resetMissingClock;
enableRx;

END;      { transmitFrame }

```

ReceivePacket Procedure

The procedure `receivePacket` is the primary interface routine to higher levels. It is specified as if it is synchronously called by the user. Note that in many implementations, the lower-level `ReceiveLinkMgmt` function would be invoked by an interrupt routine.

```

PROCEDURE receivePacket ( VAR dstParam : anAddress; VAR srcParam : anAddress;
                          VAR typeParam : aLAPtype; VAR dataParam : aDataField;
                          VAR dataLength : INTEGER);

VAR status : ReceiveStatus;

BEGIN

REPEAT
    status := ReceiveLinkMgmt;
    IF status = receiveOK THEN BEGIN
        WITH incomingPacket DO BEGIN
            dstParam := dstAddr;
            srcParam := srcAddr;
            typeParam := lapType;
            dataParam := dataField;
        END; { WITH }
        dataLength := incomingLength;
    END; { IF }
UNTIL status = receiveOK

END;      { receivePacket }

```

ReceiveLinkMgmt Function

The function `ReceiveLinkMgmt` implements the receiver side of the Link Access Protocol; it would typically be called from an interrupt routine rather than `receivePacket`.

```

FUNCTION ReceiveLinkMgmt : ReceiveStatus;

VAR

```

```
status : ReceiveStatus;
CTSframe, ACKframe : aFrame;

BEGIN

status := Receiving;
WHILE status = Receiving DO
CASE receiveFrame OF
badframeCRC, badframeSize, badframeType, underrunError, overrunError:
status := frameError;

lapENQframe : IF fAdrValid THEN BEGIN
ACKframe.dstAddr := incomingPacket.srcAddr;
ACKframe.srcAddr := MyAddress;
ACKframe.lapType := lapACK;
transmitFrame (ACKframe,3);
status := nullReceive;
END { IF }
ELSE BEGIN
fAdrInUse := TRUE;
status := nullReceive;
END; { ELSE }

lapRTSframe : IF fAdrValid THEN BEGIN
CTSframe.dstAddr := incomingPacket.srcAddr;
CTSframe.srcAddr := MyAddress;
CTSframe.lapType := lapCTS;
transmitFrame (CTSframe,3);
END { IF }
ELSE BEGIN
fAdrInUse := TRUE;
status := nullReceive;
END; { ELSE }

lapDATAframe : IF fAdrValid THEN status := receiveOK
ELSE BEGIN
fAdrInUse := TRUE;
status := nullReceive;
END; { ELSE }

noFrame: status := nullReceive;

END; { CASE }
ReceiveLinkMgmt := status;

END; { ReceiveLinkMgmt }
```

ReceiveFrame Function

The function receiveFrame is responsible for interacting with the hardware.

```
FUNCTION receiveFrame : FrameStatus;

VAR rcvtimer : REAL;

BEGIN
```



```

{ Provide timeout for idle line }
rcvtimer := RealTime + maxIDGtime;
REPEAT UNTIL CarrierSense OR (RealTime > rcvtimer);
IF NOT CarrierSense THEN BEGIN
    receiveFrame := noFrame;
    EXIT (receiveFrame);
END; { IF }

{ Line is not idle, check if frame is for us }
rcvtimer := RealTime + maxIFGtime;
REPEAT UNTIL RxCharAvail OR (RealTime > rcvtimer);
IF RxCharAvail THEN BEGIN { receive frame }
    error := FALSE; framedone := FALSE; incomingLength := 0;
    REPEAT
        rcvtimer := RealTime + 1.5 * byteTime;
        REPEAT UNTIL RxCharAvail OR (RealTime > rcvtimer);
        IF RxCharAvail THEN BEGIN
            IF OverRun THEN BEGIN
                receiveFrame := overrunError;
                error := TRUE;
            END { IF OverRun }

            ELSE IF incomingLength < maxFrameSize THEN BEGIN
                incomingLength := incomingLength + 1;
                incomingPacket.rawData [incomingLength] := rxDATA;
            END { ELSE IF }

            ELSE BEGIN
                receiveFrame := badframeSize;
                error := TRUE;
            END; { ELSE }

            IF EndOfFrame THEN
                IF CRCok THEN BEGIN
                    incomingLength := incomingLength - 2; { account for CRC }
                    IF incomingLength < minFrameSize THEN BEGIN
                        receiveFrame := badframeSize;
                        error := TRUE;
                    END { IF incomingLength ... }
                    ELSE framedone := TRUE;
                END { IF CRCok }
                ELSE BEGIN { bad CRC }
                    receiveFrame := badframeCRC;
                    error := TRUE;
                END;
            END { IF RxCharAvail }

            ELSE BEGIN { RealTime > rcvtimer }
                receiveFrame := underrunError;
                error := TRUE;
            END { ELSE }

        UNTIL framedone OR error;

        { Check on validity of the frame }
        IF framedone THEN
            IF fAdrValid THEN
                { if our address is valid, check on actual type }

```

```
IF incomingPacket.lapType < $80 THEN receiveFrame := lapDATAframe;
ELSE CASE incomingPacket.lapType OF
    lapENQ : receiveFrame := lapENQframe;

    lapACK : BEGIN
        receiveFrame := lapACKframe;
        fAdrInUse := TRUE;
    END; { lapACK }

    lapRTS : receiveFrame := lapRTSframe;

    lapCTS : IF fCTSexpected THEN receiveFrame := lapCTSframe
    ELSE BEGIN
        fAdrInUse := TRUE;
        receiveFrame := badframeType
    END; { ELSE }

    OTHERWISE receiveFrame := badframeType;
END { CASE }
ELSE IF incomingPacket.rawData[1] <> $FF THEN BEGIN
    fAdrInUse := TRUE; { We received something we didn't expect }
    receiveFrame := noFrame;
END { ELSE IF }
END { IF RxCharAvail }
ELSE receiveFrame := noFrame; { no CharAvail }

resetRx; resetMissingClock

END; { receiveFrame }
```

Miscellaneous Functions

The following low-level routines are referenced by the above code:

```
FUNCTION bitCount (bits : bitVector) : INTEGER;
VAR i, sum : INTEGER;
BEGIN
    sum := 0;
    FOR i := 0 TO 7 DO sum := sum + bits [i];
    bitCount := sum
END; { bitCount }

FUNCTION min (val1, val2 : INTEGER) : INTEGER;
BEGIN
    IF val1 < val2 THEN min := val1
    ELSE min := val2
END; { min }

FUNCTION max (val1, val2 : INTEGER) : INTEGER;
BEGIN
    IF val1 > val2 THEN max:= val1
    ELSE max:= val2
END; { max }

FUNCTION Random (maxval : INTEGER) : INTEGER;
BEGIN
    { this function is implemented as any "good" pseudo-random number generator
```

which produces a result in the range 0 .. maxval-1 }
END;

SCC Implementation

One of the integrated circuits most commonly used in the implementation of ALAP is the Zilog 8530 Serial Communications Controller (SCC). This section explains how the hardware interface routines declared above could be implemented with that device. This discussion should not be construed as in any way implying that the SCC must be used in the implementation of ALAP. Many other devices can be employed to effectively implement ALAP. Note that all of the registers and bit names below are those used by Zilog in its SCC documentation.

CarrierSense indicates that the hardware is sensing a frame on the bus; this corresponds to the complement of the SYNC/HUNT bit in RR0.

RcvDataAvail indicates that a data byte is available; this corresponds to the Rx Character Available bit in RR0.

rxDATA is the next byte from the bus; it is RR8.

EndOfFrame indicates that a valid closing FLAG has been detected; this is the End of Frame bit in RR1.

CRCok indicates that the received frame's FCS was correct (when **EndOfFrame** is true); this is the complement of the CRC/Framing Error bit in RR1.

OverRun indicates that the code did not keep up with data reception; this is the Rx Overrun Error bit in RR1.

MissingClock indicates that the hardware has detected a missing transition on the bus; this is the One Clock Missing bit in RR10.

setAddress is a procedure which sets the hardware to receive packets whose destination address matches **MyAddress**; this simply sets WR6 in the SCC.

enableTxDrivers and **disableTxDrivers** control the operation of the RS-422 drivers. This would generally be controlled by one of the SCC's output bits (on the Macintosh, it's the RTS bit in WR5).

enableTx and **disableTx** control the operation of the transmitter; this is done by means of the Tx Enable bit in WR5.

txFlag causes the transmission of a FLAG. This happens automatically at frame opening when Tx Enable is set; however code must delay long enough to cause the extra FLAG. The trailing FLAG is generated automatically at frame end as part of the Tx Underrun processing.

txDATA causes the transmission of a data byte; this is WR8.

txFCS causes the transmission of the Frame Check Sequence. This is caused automatically by letting Tx Underrun occur.

txONES causes 12+ ones (1's). This is done by disabling the SCC transmitter (setting TX Enable to zero) while leaving the RS-422 drivers on and delaying.

resetRx, enableRx and disableRx control the receiver; this is done by means of the Rx Enable bit in WR3. resetRx should also flush the receive FIFO.

resetMissingClock causes the MissingClock indication to be cleared; this is done by a Reset Missing Clock command via WR14.


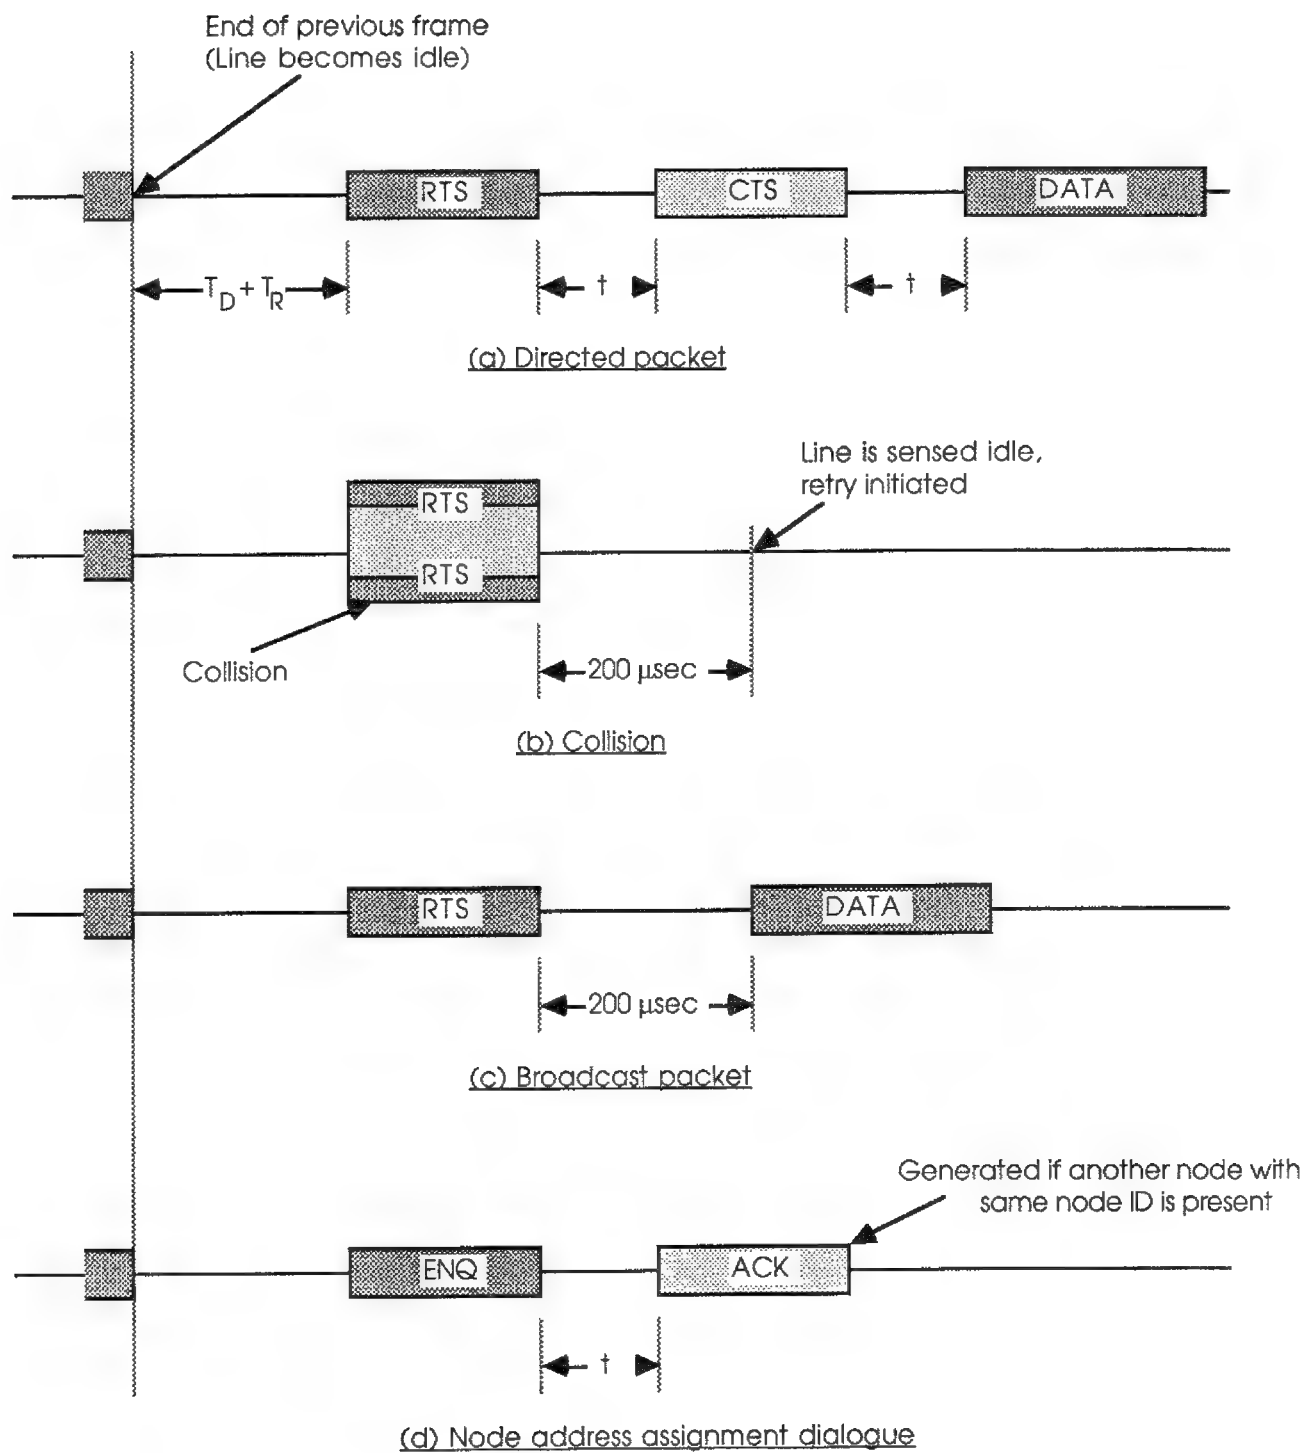


Figure III-1. ALAP Frame Format



t -- Inter-Frame Gap (less than 200 μ secs)

T_D -- Inter-Dialogue Gap (greater than 400 μ secs)

T_R -- randomly generated time interval

Figure III-2. ALAP Timing Diagrams

IV. Datagram Delivery Protocol (DDP)

About Datagram Delivery Protocol

While ALAP provides a *best effort* node-to-node delivery of packets on a single AppleTalk network, the Datagram Delivery Protocol (DDP) is designed to extend this mechanism to the socket-to-socket delivery of datagrams over an AppleTalk internet (see Figure IV-1). An AppleTalk internet consists of one or more AppleTalk networks connected by intelligent nodes referred to as *bridges* or *internet routers*. (Bridges should not be confused with *gateways*, which are nodes that separate and manage communication between different types of networks.)

Bridges are packet-forwarding agents. Packets can thus be sent between any two nodes of an internet by using a *store and forward* process through a series of internet routers. By combining this facility with DDP, packets can be sent between any two sockets on an internet.

A bridge will often consist of a single node connected to two AppleTalk networks, or it might consist of two nodes, each of which is connected to the other through a communication channel. As far as the protocol architecture is concerned, the channel between the two halves of a bridge could take a variety of forms: a leased or dial-up line, another network (for example, a wide-area packet-switched or circuit-switched public network), or a higher-speed broadband or baseband local-area network used as a *backbone*.

Sockets and Socket Identification

Sockets are logical entities within the nodes connected to an AppleTalk internet. Sockets are owned by *socket clients*, which are typically processes (or functions in processes), implemented in software in the node. A socket client can send and receive *datagrams* (this term is used to signify packets of data carried by DDP between the sockets of an internet) only through sockets that it owns. Each socket within a given node is identified by an eight-bit socket number. Socket numbers are treated as unsigned integers. There can be at most 254 different socket numbers in a given node (the values 0 and 255 are reserved and cannot be used to identify sockets).

Sockets are classified into two groups: *statically-assigned* and *dynamically-assigned*. Statically-assigned sockets (SAS) have socket numbers in the range 1 through 127; these sockets are reserved for use by clients such as the AppleTalk core protocols (for example, ATP, NBP, RTMP) and for low-level built-in network services such as echoers. Socket numbers decimal 1 through 63 are specifically reserved for use by Apple. Socket numbers decimal 64 through 127 are available for unrestricted experimental use. Use of these experimental SASs is not recommended for commercial products, since there is no mechanism for eliminating conflicting usage of the same socket(s) by different developers (see the section "Use of Name-Binding Protocol with Sockets"). See Appendix A for a summary of socket number usage.

Socket numbers decimal 128 through 254 are assigned dynamically by DDP upon request from clients in that node; sockets of this type are said to be dynamically assigned (DAS).

No two sockets in the same node can have the same socket number. Thus, the socket number taken together with the ALAP node ID provides an unique identifier for any socket on a single AppleTalk network. This is called the socket's *AppleTalk address*.

A *network number* is a 16-bit number that uniquely identifies a network in an internet. By assigning a network number to each AppleTalk network in an internet, it is possible to uniquely identify any socket on the internet. The *internet address* of a socket consists of its socket number, the node ID (of the node in which the socket is located) and the network number (of the network on which the node is located). Thus the source and destination sockets of a datagram can be fully specified by their internet addresses.

The network number 0 is reserved to mean unknown; by default it specifies the local network to which the node is connected. Packets whose destination network number is zero are addressed to a node on the local network. This allows systems consisting of a single AppleTalk network to operate without the need for an explicit network number. The value 65,535 (all bits set to one) is reserved for future use. Although we do not expect our clients to build such a large internet, the two-byte network number theoretically allows proper operation on an internet with up to 65,534 AppleTalk networks.

DDP Protocol Type Field

The AppleTalk architecture allows the implementation of a large number (up to 255) of parallel protocols that are clients of DDP (and hence are located above DDP). It is important to realize that socket numbers are not associated with a particular protocol type and should not be used to demultiplex among parallel protocols at the transport level. Instead, for this purpose a one-byte *DDP protocol type field* is provided in the DDP header. See Appendix A for a summary of the usage of the DDP protocol type field.

Socket Listeners

Socket clients provide code, referred to as the *socket listener*, that *receives* datagrams addressed to that socket. The specific implementation of a socket listener is node-dependent. For efficiency, the socket listener should be able to receive datagrams asynchronously through either an interrupt mechanism or an I/O request completion routine.

The code that implements the DDP in the node must contain a data structure called a *sockets table* to maintain an appropriate descriptor of each active socket's listener.

DDP Interface

As shown in figure IV-2, the DDP interface is the boundary at which the socket client can issue calls to and obtain responses from the DDP implementation module in the node. The DDP Implementation Module supports four calls, described below.

Opening a Statically-Assigned Socket

The caller specifies the socket number and the socket listener for that socket. The call returns with a result code:

- success: socket activated
- error: various cases: socket already active, not a statically-assigned socket (outside the permissible range), socket table full

Opening a Dynamically-Assigned Socket

This is similar to the preceding except that the caller does not specify the socket number. The call returns a result code and, if successful, the activated socket's number. The result code takes the following possible values:

- success: socket activated
- error: various cases: socket table full, all dynamic sockets busy

Closing a Socket

This request specifies the number of the socket to be deactivated. If the socket is currently active, it is removed from the sockets table. The result code has the following possible values:

- success: socket deactivated
- error: no such socket

Sending a Datagram

The request specifies the number of the source socket, the internet address of the destination socket and the DDP protocol type field value. Also the length and location of the data part of the datagram are provided in the request. Since DDP includes an optional software checksum in internet datagrams, the requester must specify whether or not this checksum is to be generated. The result code has the following possible values:

- success: datagram sent
- error: sending socket not active or valid, datagram too long

In addition to these four calls, a socket listener must provide a mechanism for the reception of datagrams.

Datagram Reception by the Socket Listener

We do not specify this function since it is dependent on the implementation in the node. Some mechanism is needed to deliver datagrams within the node to the destination socket's listener. The DDP package should attempt this only if the destination socket is currently active. *The DDP must discard datagrams addressed to inactive sockets.* Furthermore, internet datagrams received with an invalid DDP checksum must be-discarded.

DDP Internal Algorithm

DDP is a simple, best-effort protocol for delivery of datagrams. As such there is no mechanism for recovery from packet loss or error situations.

Within the DDP implementation module, the primary function is to form the DDP header on the basis of the destination address, and then to pass the packet on to the appropriate ALAP. Similarly, for packets received from ALAP, DDP must examine the datagram's destination address in the DDP header and route the datagram accordingly. Details of this operation depend on whether the node is a bridge; this is discussed more fully in the section *DDP Routing Algorithm*.

DDP Packet Format

A datagram consists of the DDP header followed immediately by the data.

There is a 10-bit datagram length field in the DDP header (see Figures IV-3 and IV-4). The value in this field is the length in bytes of the datagram starting with the first byte of the DDP header and including all bytes up to the last byte of the data part of the datagram. Upon receiving a datagram, the receiving node's DDP implementation must reject all datagrams whose indicated length is not equal to the actual received length. The maximum length of the data component of a datagram is 586 bytes; longer datagrams must be rejected.

Short and Long Form Headers

In addition, the DDP header contains the source and destination socket addresses and the DDP protocol type. Each of these addresses can be specified as a four-byte internet address. However, for packets whose source and destination sockets are on the same network, the network number fields are unnecessary; likewise, for such datagrams the source and destination node IDs are already found in the ALAP header, and would thus be redundant in the DDP header. For these reasons, DDP uses two types of header: a short form and an extended (or long) form.

DDP uses the value of the ALAP protocol type field to determine if the packet has a short or an extended DDP header. The ALAP protocol type field value is 1 for the short form, and 2 for the extended form.

The short form version of a DDP datagram is shown in Figure IV-3. The datagram header is five bytes long. The first two bytes of the header contain the datagram length, with the more significant byte first. The upper 6 bits of this byte are not significant and should be set to zero. The datagram length field is followed by a one-byte destination socket number, a one-byte source socket number, and a one-byte DDP protocol type field. Such short header datagrams are sent if the source and destination sockets are on the same network.

The extended form datagram is shown in Figure IV-4. The extended form DDP header is thirteen bytes long. It contains the full internet addresses of the source and destination sockets, as well as the datagram length and DDP type fields. For such packets, there is a 6-bit *hop count* field (described below) in the most significant bits of the first byte of the DDP header. Datagrams exchanged between sockets on different AppleTalk networks of an internet must use this form of header. In addition, the extended header includes a two-byte (16-bit) DDP checksum field. If the sending client so desires, the source node's DDP implementation calculates and inserts a software-generated checksum into this field;

otherwise, the sending node sets this field to 0. The datagram's destination node recomputes this checksum and rejects the datagram if the received and computed values do not agree. All two-byte fields are specified high byte first.

The DDP checksum has been provided to allow the detection of errors caused by faulty operation (for example, memory and data bus errors) within bridges on the internet. Implementors of DDP should treat it as an optional feature.

DDP Checksum Computation

The 16-bit DDP checksum is computed as follows:

```
CkSum := 0 ;
FOR EACH datagram byte starting with the byte immediately following the CkSum
REPEAT the following algorithm:
    CkSum := CkSum + byte; (unsigned addition)
    Rotate CkSum left one bit, rotating the most significant bit into the least
    significant bit;
IF, at the end, CkSum = 0 THEN
    CkSum := hex FFFF (all ones).
```

Reception of a datagram with CkSum = 0 implies that it is not checksummed.

Hop Counts

For datagrams that are exchanged between sockets on two different AppleTalk networks in an internet, a provision is made to limit the maximum number of internet routers the datagram can visit. This is done by including, in such internet datagrams, a hop count field.

The source node of the datagram sets this field to zero before sending the datagram. Each internet router increments this field by one. A bridge receiving a datagram with a hop count value of 15 should not forward it to another bridge; if such a datagram's destination is on a local network directly connected to the bridge, then the bridge should send it to that destination, otherwise the datagram should be discarded by the bridge. This provision is made to filter out of the internet packets that might be circulating in closed routes. Such closed routes (loops) are a transient situation that can occur for short periods of time while the routing tables are being updated by the Routing Table Maintenance Protocol (RTMP). Non-bridge nodes ignore this field.

The upper two bits of the hop count currently are not used by DDP, but are reserved for future use (such as the extension of the maximum value of the hop count beyond the currently allowed value of 15).

DDP Routing Algorithm

A datagram is conveyed from its source to its destination socket over the internet through the bridges. The DDP implementation in the source node examines the destination network number of the datagram and determines whether the destination is on the local network or not. If it is, the short form DDP header is adequate, and the ALAP layer is called to send the packet to its destination node. If, however, the destination is not on the local network, DDP builds the extended header and calls ALAP to send the packet to a bridge (if there is more than one such bridge, any one will do) on the local network. The bridges examine the

destination network number of the datagram, and use routing tables to forward the datagram to subsequent bridges (through the ALAPs of appropriate intervening local networks to which the bridge is connected) to get it to a bridge connected to the destination network. There the datagram is sent to its destination node through the local network's link level protocol.

Each node on an AppleTalk maintains as an internally stored value the network number of the local network to which it is attached. The DDP implementation in a datagram's source node determines if the destination network is the local network or not by comparing the destination network number to the internally stored value of the local network's number. If the two numbers are the same then a short DDP header is built for the datagram. An important special case arises when the internally stored value is zero (unknown). In this case, if the datagram's destination network number is non-zero, then DDP should assume that the packet is intended for a node on the local network. However, it must, in this case, build an extended DDP header and call ALAP to send the packet to the specified destination node *on the local net*. The extended DDP header enables the receiving node to throw away the packet if it determines it was not the intended recipient (the destination network of the packet is not equal to its internally stored local network number).

Routing tables are maintained by bridges by using the Routing Table Maintenance Protocol (RTMP) discussed in a separate chapter. The routing tables indicate, for each network number in the internet, the node ID (on the appropriate local network) of the next bridge on the proper route.

Non bridge nodes do not need to maintain these tables. Such nodes need only two pieces of information: the network number of the local network (THIS-NET), and the node ID of any bridge (A-BRIDGE) on the local network. This can be done by implementing a simple subset of RTMP, called the *RTMP Stub*, in each such node. For nodes on systems consisting of a single network, the values of THIS-NET and A-BRIDGE will be zero (*unknown*).

The internal routing algorithm of the DDP implementation module for a non-bridge node is given below. The sending client issues the send datagram call, specifying the destination socket's internet address.

```
IF (destination network number = 0) OR (destination network number = THIS-NET) THEN
    BEGIN
        build the short-form DDP header;
        call ALAP to send it to the destination node
    END
ELSE
    BEGIN
        build the extended-form DDP header;
        IF THIS-NET = 0 THEN call ALAP to send the packet to the destination node
        ELSE IF A-BRIDGE <> 0 THEN call ALAP to send the packet to A-BRIDGE
        ELSE return an error (no bridge available)
    END;
```

For packets received from ALAP by non bridge nodes, the routing function is simply one of delivering the packet to the destination socket in the node. DDP must first verify that the destination network number in an extended DDP header is in fact equivalent to that node's internally stored value of THIS-NET, and ignore the packet if not. (For a precise definition of this equivalence see the discussion at the end of this chapter.) It is also advisable for such nodes to verify that the destination node ID in an extended DDP header matches the station's node ID (or is equal to the broadcast address, \$FF).

In internet routers, the routing algorithm is somewhat more complex; details are provided in the chapter on RTMP.

Use of Name-Binding Protocol in Conjunction With Sockets

Developers of products for AppleTalk should not use statically assigned sockets except for purely experimental purposes. This restriction is imposed to avoid the conflicting use of the same statically assigned socket by different developers. Such conflicts are difficult to avoid in the absence of a central administering body.

Developers should instead use the name-binding technique to allow workstations to discover their server/service access point addresses. Thus developers would identify their server service by a unique name. Workstations would then use NBP to bind an address to this name (for details, see the section on NBP). Once the client process has determined the proper destination socket number, it can then proceed to transmit packets to that socket.

Clearly, this requires that developers must implement NBP in their servers. While not significant for larger complex servers, this could pose a significant problem for smaller, memory-bound cases. In fact, NBP has been so designed that a subset is all that is needed for such servers. This subset simply responds to *lookup* packets received over the network. Since the names table of such a server will contain only a single name, this NBP subset need not implement functions such as names table management, etc.

Network Number Equivalence

The use of network number zero to indicate *unknown* introduces some complications for DDP clients. A DDP client may wish to compare two network numbers to determine if they are equivalent. For instance a request may be sent to a node on network 7 and a response received from one on network zero. Was the response received from the network to which the request was sent? For this reason, we must somewhat arbitrarily define when two network numbers match (i.e. are equivalent). The rule to use is *zero matches anything*. Thus network A is equivalent to network B if $A=B$ or $A=0$ or $B=0$. All DDP clients should use this definition of network equivalence.

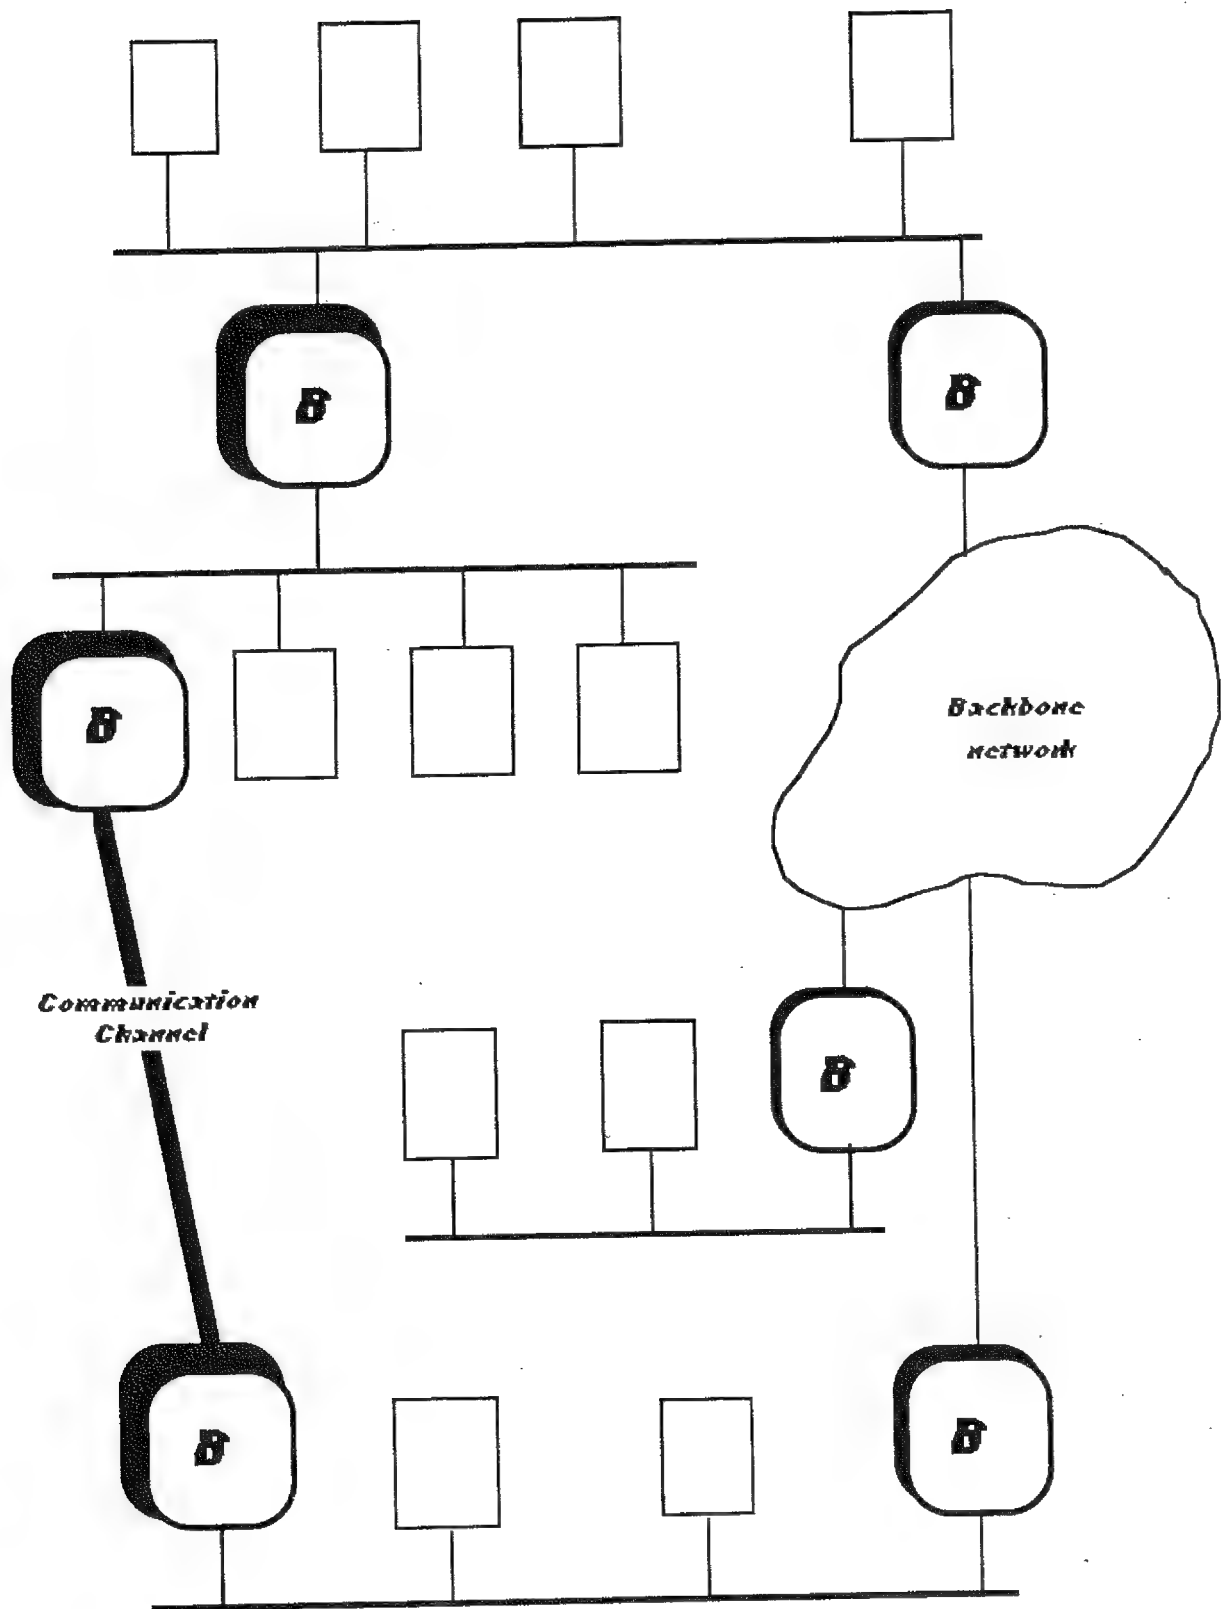


Figure IV-1: AppleTalk internetwork and bridges/internet routers (B)

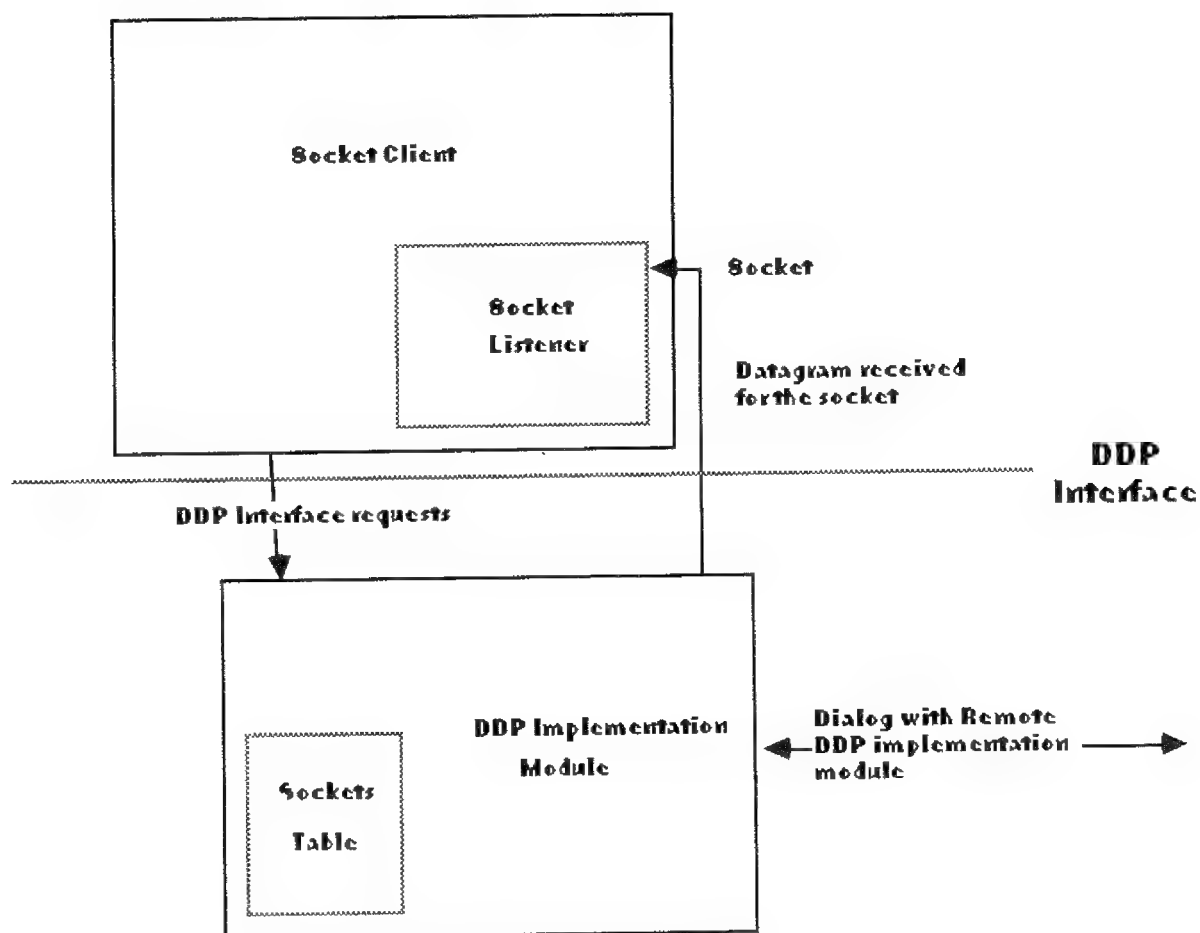


Figure IV-2: Socket Terminology

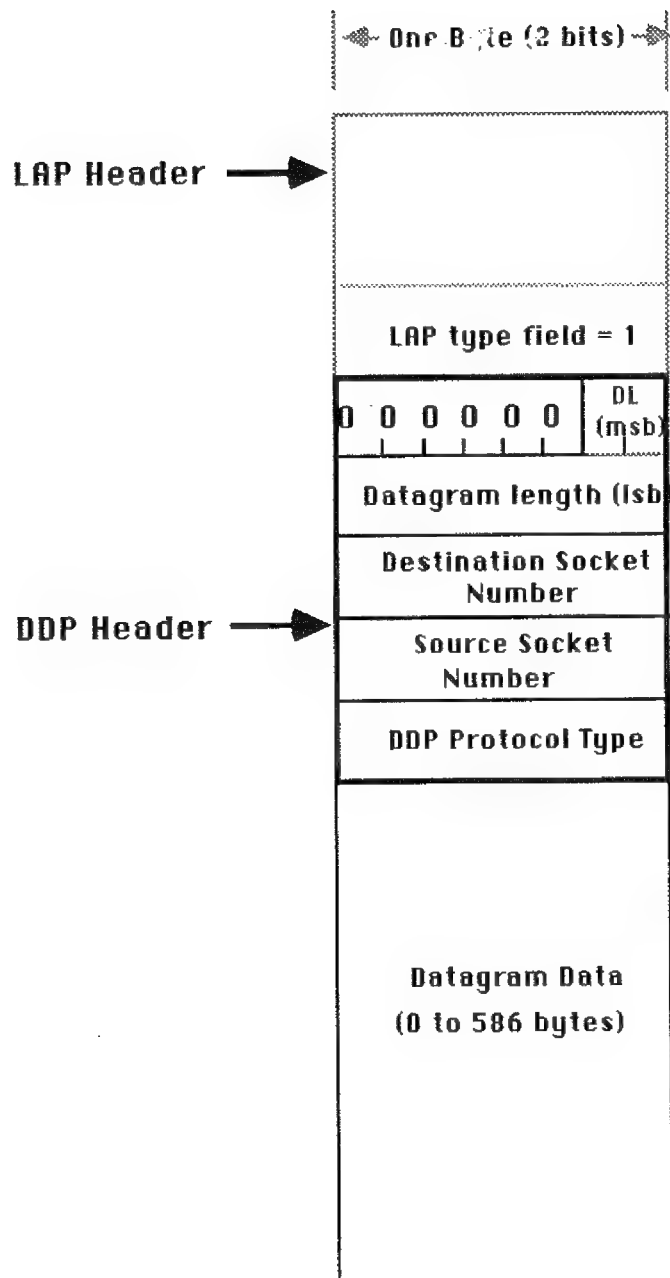


Figure ID-3: DDP Packet Format

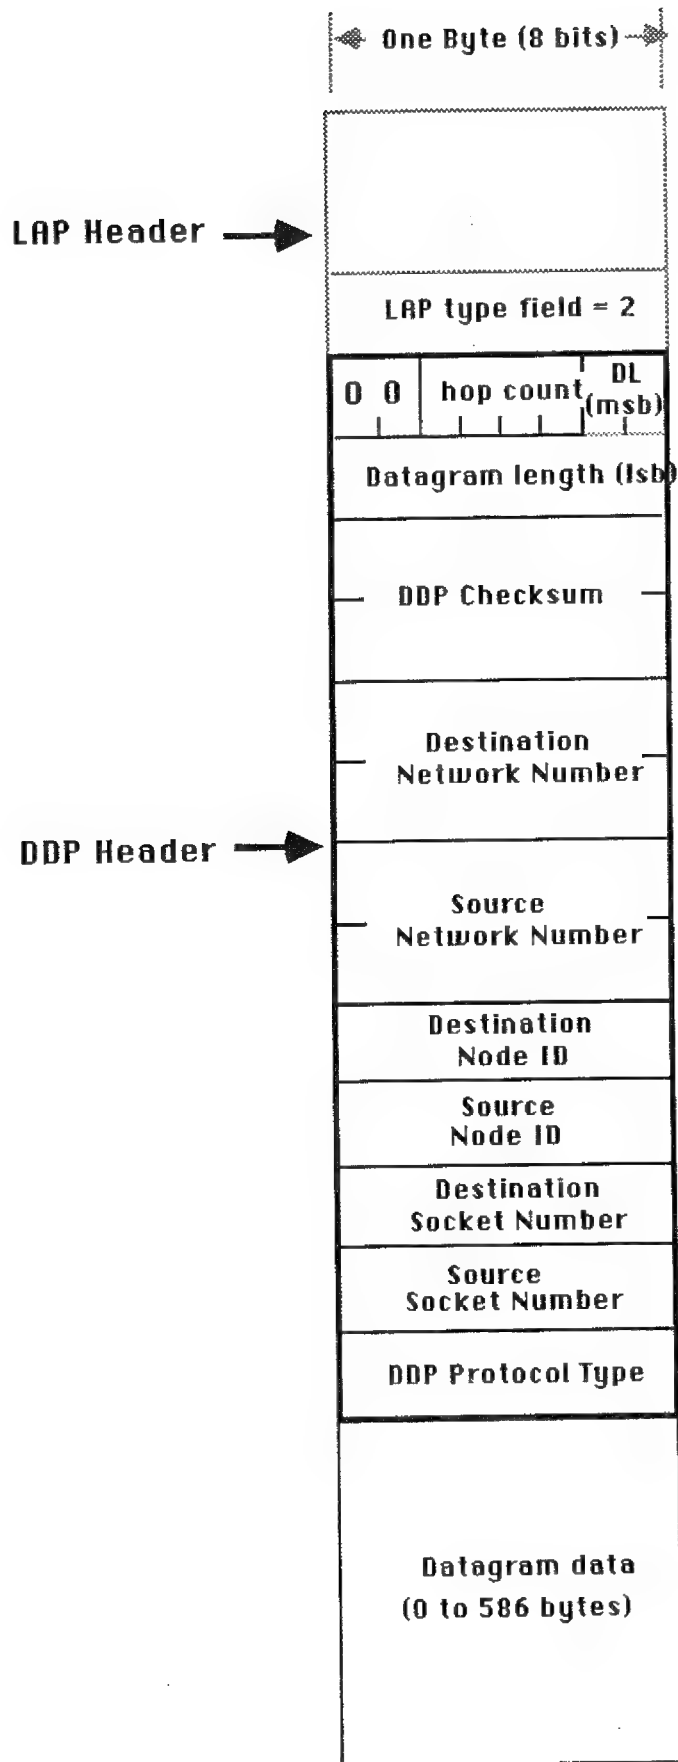


Figure IV-4: DDP Packet Format

V. Routing Table Maintenance Protocol (RTMP)

About Routing Table Maintenance Protocol

The Datagram Delivery Protocol introduced the concept of bridges or internet routers (IR) as the mechanism by which datagrams are forwarded/routed from any source socket to any destination socket on the internet. The Routing Table Maintenance Protocol (RTMP) is used by bridges to establish and maintain the *routing tables* that are central to this forwarding process.

Bridges

Bridges are the key components in extending the datagram delivery mechanism to an internet setting. It is useful to distinguish the ways in which bridges can be used to build an internet; Figure V-1 illustrates three principal situations.

Local Bridges

Configuration A shows a bridge used to interconnect several AppleTalk networks that are in close proximity; such a bridge can be said to be a *local bridge*. Such local IRs are connected directly to each of the AppleTalk networks they serve to bridge. They are useful in allowing the construction, within the same building, of an AppleTalk internet with more than 32 nodes.

Half Bridges

Configuration B illustrates the use of two bridges interconnected by a long distance communication link. Each bridge is directly connected to an AppleTalk network. The combination of the two bridges and the intervening link serves in effect as a bridging unit between the AppleTalk systems. Each bridge in this unit will be referred to as a half bridge. The intervening link can of course be made up of several devices (such as modems), and other networks (such as public data networks). The primary use of half bridges is to interconnect remote AppleTalk systems. An important characteristic of such bridges is that the throughput is generally lower than in the case of local bridges, due to the addition of the communication link. Furthermore, these communication links can often be less reliable than the local networks of the internet.

Backbone Bridges

Although these bridges might be placed in one of the previous categories, they present an important set of properties that make it appropriate to single them out. These bridges are used to interconnect several AppleTalk networks through a backbone network. Each bridge could be a local bridge connected on one side to an AppleTalk network, and to the backbone network on the other. This is illustrated as configuration C. Another manner of connecting a backbone bridge to the backbone network might be through a long-distance communication link.

Bridge Model

We model a bridge as a device with several hardware ports, referred to as *bridge ports* (see Figure V-2). A bridge port can be connected in any of the three ways described above (local to an AppleTalk, half to a communications link, or backbone to a high-speed net), or even as a combination of all three. In our model a bridge can have any number of ports, which are numbered starting with the number 1.

Each bridge port has associated with it a *port descriptor*. This consists of four fields: (1) a flag as to whether the port is connected to an AppleTalk network or not, (2) the port number, (3) the port node ID, (the bridge's node ID corresponding to that port), and (4) the port network number, (the network number of the network to which the port is connected).

The values of these four fields are obvious for a port that is directly connected to a local AppleTalk network. For a port connected to one end of a communication link (half bridge case), the port node ID and port network number are meaningless. For a port connected to a backbone network, the port network number is meaningless, while the port node ID is the appropriate node ID of the bridge on the backbone network. In this latter case, provision must be made for this field to be of any size (possibly variable length) depending on the nature of the backbone network.

It is important to understand that the AppleTalk node ID of a local bridge is different for each of its ports. In other words, for each AppleTalk network to which it is directly connected, the bridge acquires a different node ID.

The bridge internals include a suitable data link process (ALAP or other) for each port, a (DDP) router, the routing table, and the routing table maintenance process (RTMP) implemented on a statically-assigned socket known as the RTMP socket (socket number = 1). The router accepts incoming datagrams from the LAPs and then reroutes them out through the appropriate port depending on their destination network number. This routing decision is made by the router by consulting the routing table. The routing table maintenance process receives RTMP packets, from other bridges, through the RTMP socket and uses these to maintain/update the routing table. (Bridges additionally consist of a Zone Information Protocol process and an NBP routing process; these are discussed in the ZIP chapter).

Internet Topologies

RTMP allows internets to consist of AppleTalk networks interconnected via bridges in any arbitrary topology. The only strict limitation imposed on an AppleTalk internet is that, for each bridge, no two of its ports be on the same network. In addition, nodes on any network that is more than sixteen hops away (via the shortest path) from any other network will not be able to communicate with nodes on that network.

Routing Tables

All bridges maintain complete routing tables that allow them to determine how to forward a datagram on the basis of its destination network number. RTMP allows bridges to periodically exchange their routing tables; a bridge receiving the routing table of another bridge compares it with its own table, updating its own to record the shortest path for each destination network. This exchange process allows the bridges to respond to changes in

the connectivity of the internet (such as when a bridge goes down or a new bridge is installed).

A routing table that has stabilized to all changes will consist of one entry for each reachable network in the internet. Each entry provides the number of the port through which packets for that network must be forwarded by the bridge, the node ID of the next IR/bridge, and a measure of the *distance* to the destination network. The entry in the routing table corresponds to the shortest path (known to that bridge) for the corresponding destination network. The distance corresponding to a network to which the bridge is directly connected is always zero.

The distance to a network is measured in terms of *hops*, each hop representing one IR/bridge that the packet will encounter in its path from the current bridge to the destination network. This simple measure of distance is adequate for an RTMP that adapts to changes in the connectivity of the network.

(Other modified measures could reflect the speeds/capacities of the intervening links and thus try to find a minimum time path. Yet another enhancement might use current traffic conditions on a particular path to modify its contribution to distance. We have, for sake of simplicity, chosen to use the hop-count measure. The basic algorithm remains unchanged if more complex measures are used.)

Each table entry has associated with it an *entry state value*. This is a variable that takes on one of three values: good, suspect, bad. The significance of the entry state will become clear when the table maintenance mechanism is discussed in more detail.

Figure V-3 illustrates a typical routing table for a bridge with three ports in an internet consisting of seven networks. The corresponding port descriptors are shown in the figure as well.

Routing Table Maintenance

Since bridges have no knowledge of the topology or connectivity of the internet, RTMP must provide the mechanism for constructing their routing tables, and maintaining them in the face of changes in the internet (bridges coming up or going down).

When a bridge is switched on, it initializes its table by examining the port descriptor of each of its ports. Any AppleTalk port with a non-zero port network number signals that the bridge is directly connected to that network. Thus an entry is created in the table for that network number, with a distance of zero and with that port's number in the appropriate part of the entry. This initial table is called the *routing seed* of the bridge.

Every bridge must periodically broadcast one or more *RTMP data packets* through each of its ports, addressed as datagrams to the RTMP socket. Thus, every bridge's RTMP will periodically receive RTMP data packets from every bridge on its directly connected networks, backbones or communication links. The RTMP data packets (see Figure V-4) carry the port node ID and port network number of the bridge port through which the RTMP data packet was sent, as well as the <network number, distance> pairs (called *routing tuples*) of the entries in the sending bridge's routing table. Using these RTMP data packets, RTMP adds to or modifies its own routing table.

The basic idea is that if an RTMP data packet (received by the bridge) contains a routing tuple for a network not in the bridge's table, then an entry is added for that network number

with a distance one larger than the tuple's distance. In essence, the RTMP data packet indicates that a route exists to that network through the packet's sender.

Likewise, if an RTMP data packet indicates a shorter path to a particular network than the one in the bridge's routing table (that is, if the tuple distance plus one is less than the table entry's distance), then the corresponding entry must be modified to indicate the RTMP data packet's sender as the next IR for that network with the new distance. In fact, even if the paths are of equal length, the entry is modified (thus information remains as up to date as possible).

Clearly, this process allows for the growth and adaptation of routing tables to the addition of bridges/routes.

"Aging" Routing Table Entries

However, if bridges die or are switched off, the corresponding changes will not be discovered through the foregoing process. To respond to such changes, the entries in the routing tables must be *aged* and, in the absence of confirmation via new RTMP data packets, be declared *suspect* and later *bad*. Bad entries are eventually purged from the routing tables.

Each entry in the routing table (corresponding to a network to which the bridge is not directly connected) was obtained from an RTMP data packet received by the bridge from the next IR for that network. RTMP considers such an entry to be valid for a limited time only (called the *entry validity time*). Before starting the validity timer, the bridge goes through its routing table and changes the state of every *Good* entry to *Suspect*.. An entry must be revalidated from a new RTMP data packet before the timer expires.

Suppose, for instance, that the next IR, for a particular entry in bridge B's routing table, dies. That IR will not be sending RTMP data packets; the entry's validity timer will expire and bridge B will not have received confirmation of the entry. At that time, the entry's state is changed from *Suspect* to *Bad*. Any other RTMP data packet received with path information to the network of the entry can be used to replace the entry with the new values from that packet. If no new route is discovered, however, the bad entry will be deleted when the validity timer goes off a second time.

A more detailed discussion of this process is provided in the section "RTMP Algorithms" below.

Values for the Validity and Send-RTMP Timers

These values have a significant effect on the dynamics of the propagation of routing table adaptation in response to changes in the internet's connectivity, and on network traffic. The exact values of these parameters have been determined through experimentation with actual internets. These values are 10 seconds for the Send-RTMP timer and 20 seconds for the Validity timer.

RTMP Data Packet Format

The format of an RTMP data packet is shown in Figure V-4. Clearly, the datagram need use only the short form of the DDP header. The DDP type field is set to 1 to indicate that it is an RTMP data packet. The DDP data part of the datagram consists of three parts: the sender's network number and node ID, and the routing tuples.

Sender's Network Number

The first two bytes of the RTMP data packet's DDP data is the port network number from the port descriptor (of the port through which the packet is sent by the bridge). This field allows the receiver of the packet to determine the number of the network through which the packet was received. This is thus the number of the network to which the corresponding port of the receiver is attached. RTMP Data packets sent out ports that are not AppleTalk networks should have this field set to zero.

Sender's Node ID

This field contains the port node ID from the port descriptor (of the port through which the packet is sent by the bridge). To allow for the case of ports connected to backbone networks other than AppleTalk networks, this field must be of variable size. The first byte of the field contains the length (in bits) of the sender node's ID. This is followed by the ID itself. If the length of the ID in bits is not an exact multiple of 8, it is prefixed with enough zeros to make a complete number of bytes. The bytes of this modified ID are then packed into the ID field of the packet, most significant byte first. It is from this field that the receiver of the packet determines the ID of the bridge sending the packet.

Routing Tuples

The last part of the RTMP data packet consists of the routing tuples from the sending bridge's routing table. These <network number, distance> pairs consist of two bytes for the network number and one byte for the distance.

For internets with a large number of networks, the entire routing table may not fit in a single datagram. In that case, the tuples are distributed over as many RTMP data packets as necessary.

Assignment of Network Numbers

Network numbers are set into the port descriptors of the bridge ports, and are then dynamically propagated out through RTMP to the other nodes of each network.

Not all bridges on a particular network must have the network number set into their corresponding port descriptors. The precise requirement is that at least one bridge (called the *seed bridge*) on a network should have the network number built into its port descriptor. The other bridges could have a port network number value of zero; they will acquire the correct network number by receiving RTMP data packets sent out by the seed bridge.

An absolute requirement is that the bridges on a particular network should not have (in their port descriptors) conflicting port network numbers for that network. (The value zero does not cause a conflict).

If a bridge is not a seed bridge for a particular port, it should not send its routing table out that port, nor should it have an entry in its routing table for that port, until it acquires that port's network number. It must, however, operate correctly with regard to those ports whose network number it does know.

RTMP Initialization and Maintenance Algorithms

Initialization

A bridge upon being switched on performs the following algorithm:

```
FOR each port P connected to an AppleTalk
  IF the port network number  $\neq$  0 THEN
    create an entry for that network number:
      Entry's network no. = port network number
      Entry's distance = 0
      Entry's next IR = 0
      Entry's state = Good
      Entry's port = P;
```

This algorithm creates an entry for each directly-connected AppleTalk for which the bridge is a seed bridge.

Maintenance

The bridge is assumed to have two timers running continuously. These are the *validity timer* and the *send-RTMP timer*. The events to which the bridge's RTMP process responds are: RTMP data packet is received, the validity timer expires, the send-RTMP timer expires. The corresponding algorithms are given below:

(1) RTMP data packet is received through port P:

```
IF P is connected to an AppleTalk AND P's network number = 0 THEN
  BEGIN
    P's network number := packet's sender network number;
    IF there is an entry for this network number THEN delete it;
    Create a new entry for this network number:
      Entry's network no. = packet's sender network no.
      Entry's distance = 0
      Entry's next IR = 0
      Entry's state = Good
      Entry's port = P;
  END;

FOR each routing tuple in the data packet DO
  IF there is an entry in the table corresponding to the tuple's network number
  THEN Update-the-Entry
  ELSE Create-New-Entry;
```


General-purpose routines: Update-the-Entry, Create-New-Entry, and Replace-Entry are as follows:

Update-the-Entry

```
IF (Entry's State = Bad) AND (tuple distance < 15) THEN
  Replace-Entry
ELSE
  IF Entry's distance >= (tuple distance+1) AND (tuple distance < 15) THEN
    Replace-Entry
  ELSE
    IF Entry's next IR = data packet's sender node ID AND Entry's port = P
    THEN
      { If entry says we're forwarding to the IR who sent us this
        packet, the net's now further away than we thought }
      BEGIN
        Entry's distance := tuple distance + 1;
        IF entry's distance < 16 THEN Entry's state := Good
        ELSE Delete the entry
      END;
```

Create-New-Entry

```
Entry's network number := tuple's network number;
Replace-Entry;
```

Replace-Entry

```
Entry's distance := tuple's distance + 1;
Entry's next IR := packet sender's ID;
Entry's port number := P;
Entry's state := Good;
```

(2) Validity Timer Expires:

```
FOR each entry in the routing table DO
  CASE Entry's state OF
    Good: IF entry's distance > 0 THEN Entry's state := Suspect;
    Suspect: Entry's state := Bad;
    Bad: Delete the entry
  END;
```

(3) Send-RTMP Packet Timer Expires:

```
IF routing table is not empty THEN
  BEGIN
    Copy the network number and distance pairs of each Good or Suspect
    entry of the routing table whose distance < 15 into the routing tuple fields
    of the RTMP data packet;

    FOR each bridge port DO
      IF the port is connected to an AppleTalk AND its network no. is non-zero THEN
        BEGIN
          Packet's sender network number := port network number;
          Packet's sender node ID := port node ID;
          Call DDP to broadcast the packet to the RTMP socket;
        END
      ELSE
        IF the port is not connected to an AppleTalk THEN
          BEGIN
            Packet's sender network number := 0;
            Packet's sender node ID := port node ID;
            Call DDP to broadcast the packet through the port's LAP
            to the RTMP socket;
          END;
        END;
      END;
    END;
  END;
```

Supplements to the Algorithm for Half-Bridges

The above algorithm can be used for half-bridges, where the communications link between the half-bridges is treated as a network. In this case, the communications link will count as an additional hop. Due to the fact that this link is generally of limited bandwidth (for example, 9600 baud), sending of large routing tables between half-bridges may waste a good portion of that bandwidth, leaving little for the actual routing of packets. As an alternative, the two half-bridges plus the communications link, taken as a unit, can be considered logically to be one local bridge. In this case there is, logically, only one routing table (and Zone Information table), which is distributed between the two half-bridges. The communications link is used, not as a network, but as a medium to maintain this distributed information.

An example of this idea follows. Each half-bridge maintains a complete copy of the routing table. Whenever a half-bridge receives information (through one of its ports not connected to a communications link) that results in a change to that table, it sends only the change over the communications channel. In this way, under stable internet conditions, none of the channel will be wasted with routing information. Steps, of course, must be taken to ensure the consistency of this distributed data base, for example an acknowledgment which includes some sort of checksum of the information and the ability to re-initialize one or both sides and start from scratch.

Half-bridges must accept and process, as a minimum, RTMP Data packets as will be sent when the communications link is considered to be a network (they must also accept and process normal ZIP packets). Many half-bridges, however, will also wish to implement schemes as described above. Packets sent to implement this type of scheme should be differentiated from normal AppleTalk packets through a *non-DDP LAP type* (neither 1 nor 2). Bridges that do not understand these packets should ignore them.

Additional RTMP Services

RTMP provides a service which enables non-bridge nodes to request the number of the network they are on and the address of a bridge on that network. This service would normally be used by a node when it was first turned on. The non-bridge nodes makes the request by broadcasting, through any socket, a DDP packet of protocol type 5, addressed to the RTMP listening socket (socket 1). This packet is known as an *RTMP request*. The response, sent by any bridge receiving the request, takes the form of a normal RTMP data packet, except that it is sent directly to the requesting node and it contains no routing table information. Thus the non-bridge node can receive the response through the exact same mechanism it uses to monitor RTMP data packets (see the following section). The specific formats of the request and the response are provided in Figure V-4.

RTMP and the Zone Information Protocol (ZIP) are integrally related. ZIP uses both the routing table information and many of RTMP's significant events in its operation. ZIP also requires an additional RTMP service for the purpose of changing the zone name of a network. This service is described in detail in the ZIP chapter. Basically, ZIP must be able to request RTMP to remove an entry for a directly connected network from its routing table, and to replace that entry at a future time. The ZIP chapter details how these requests work.

RTMP and Non-Bridge Nodes

Non-bridge nodes do not need to maintain routing tables. As noted in the DDP chapter, these nodes only need to know the number of the network to which they are connected (THIS-NET) and the node ID of any bridge on that network (A-BRIDGE).

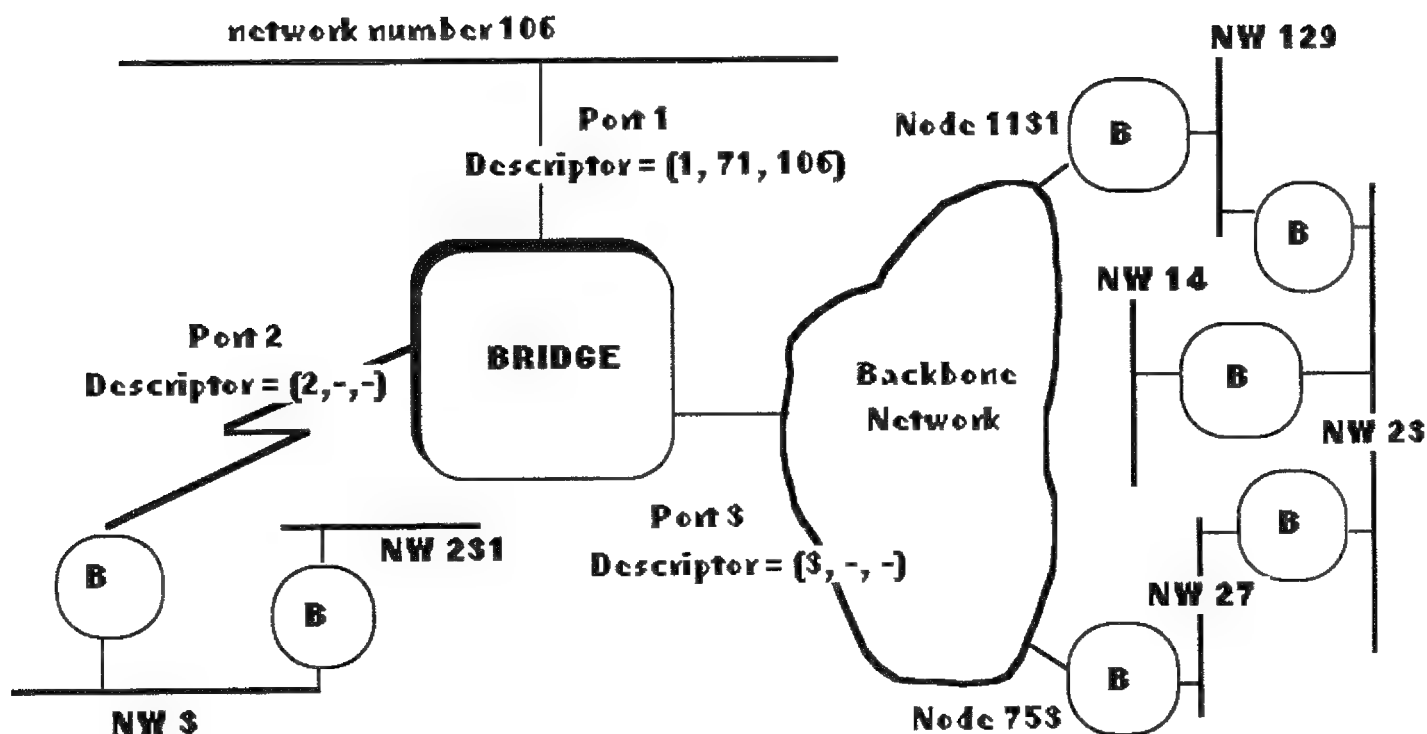
When such a node first comes up, the values of both of these variables are zero ("unknown"). These nodes can obtain the correct values dynamically by listening for RTMP data packets being sent out by the bridges on the network (if they choose to do this, they should be sure to wait long enough to get such a packet). Alternatively, they can broadcast out an RTMP request to obtain that information. In either case, to receive this information, these nodes implement a very trivial RTMP process, known as the *RTMP stub*. This process "sits" on the RTMP socket in that node, and upon receiving an RTMP data packet, copies the packet's sender network number and sender node ID fields into THIS-NET and A-BRIDGE respectively. This is done every time an RTMP data packet is received. While THIS-NET will stabilize to a constant value, A-BRIDGE may change continually (if there is more than one bridge on the network).

Additionally, such nodes must maintain a background timer for the purpose of aging this information. This is to handle the case where all the bridges connected to a given network go down and the network becomes isolated. In this case, A-BRIDGE should be reset to zero. THIS-NET, however, should not be reset. The value of this timer should be on the order of 90 seconds. Each time an RTMP data packet is received, the timer is reset. If the timer ever expires, A-BRIDGE should be aged.

Note that it is important for the RTMP stub to differentiate between RTMP data packets, broadcast by bridges, and RTMP requests, broadcast by non-bridge nodes. The RTMP stub should ignore RTMP requests. This can be done because RTMP requests have a DDP protocol type of 5, whereas RTMP data packets have a DDP protocol type of 1.

RTMP Routing Algorithm

The routing algorithm used by bridges to forward internet datagrams is given in Figure V-5. This discussion applies only to the forwarding of packets received by the bridge through one of its ports, and does not hold for packets generated within the bridge. The algorithm assumes that when a packet is received through one of the bridge ports, it is tagged with the number of the port and placed in a queue. The router takes packets off this queue, and then executes the algorithm.



Network Number	Distance	Port	Next Bridge/IR	Entry State
3	1	2	0	Good
14	3	3	1131	Good
23	2	3	753	Good
27	1	3	753	Good
106	0	1	0	Good
129	1	3	1131	Good
231	2	2	0	Good

Figure V- 3 Example of a Routing

RTMP Request

LAP Dest = \$FF
LAP Source (n)
LAP Type = 1
Length
Dest Socket = 1
Source Socket
DDP Type = 5
Command = 1

RTMP Response

LAP Dest = n
LAP Source
LAP Type = 1
Length
Dest Socket = 1
Source Socket = 1
DDP Type = 1
Sender's Network number
ID Length
Sender's ID

Node ID of the bridge sending out the RTMP packet

Routing Tuples from the sending bridge's routing table

RTMP Data Packet

LAP Dest = n
LAP Source
LAP Type = 1
Length
Dest Socket = 1
Source Socket = 1
DDP Type = 1
Sender's Network number
ID Length
Sender's ID
Network Number
Distance
...
Network Number
Distance

Figure V-4: RTMP Packet Formats

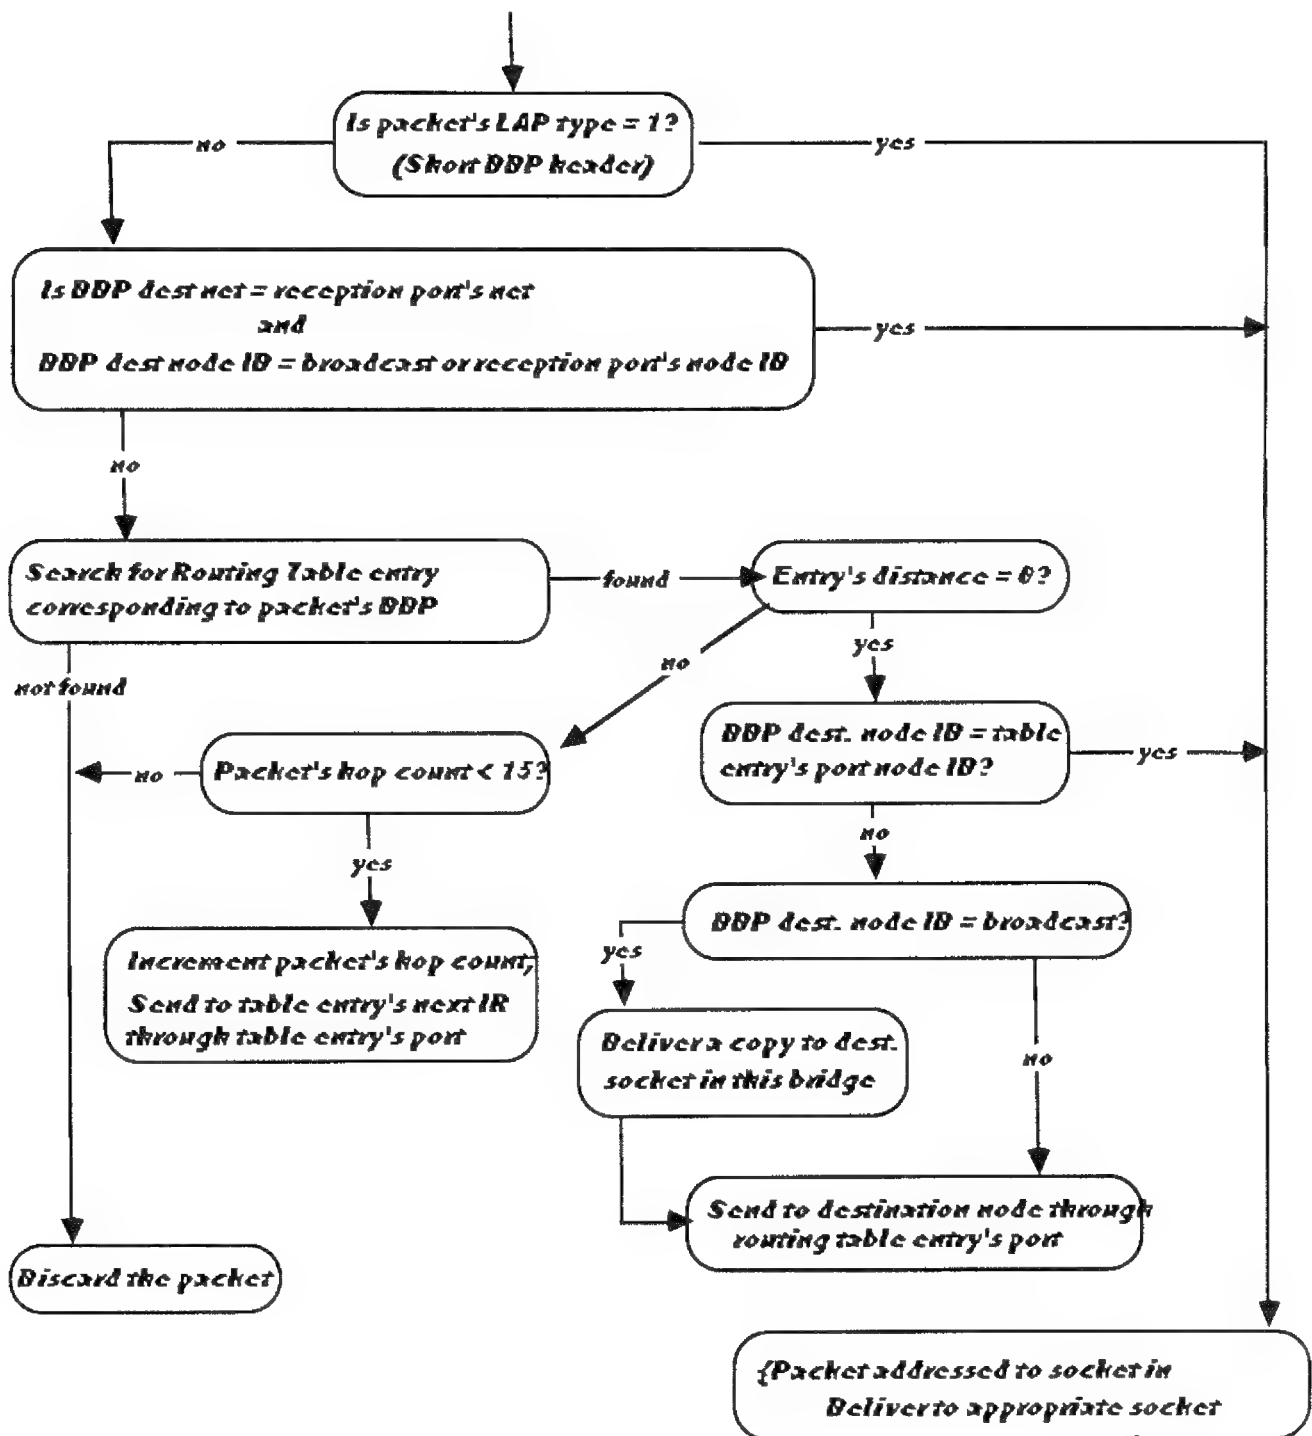


Figure V-5 : Datagram routing algorithm

VI. Name Binding Protocol (NBP)

About Name Binding Protocol

AppleTalk protocols rely on numerical identifiers, such as node IDs, socket numbers and network numbers, to provide the basic addressing capability essential for communication over a packet-switched network. However, human users may not find numbers the most convenient form of identification. Numbers are hard for them to memorize, and are easily confused and misused. Names are more easily used by the human user. Thus, if they are allowed to refer to objects by their names, the name must be converted into a network address for use by the other protocols. The Name-Binding Protocol (NBP) performs this conversion.

AppleTalk uses dynamic node ID assignment, which does not allow addresses to be configured into software for accessing network resources. The name-binding approach provides the preferable solution towards this end (see the section *Use of Name-Binding Protocol in Conjunction with Sockets in the Datagram Delivery Protocol* chapter).

Network-Visible Entities

We start by defining the concept of a *network-visible entity* (NVE). In general terms, this is any entity that can be accessed over the network. More specifically, the socket clients on an internet are its network-visible entities.

The nodes of the network are not network-visible entities; rather, any services in the nodes available for access over the network are network-visible entities. Consider, for instance, a print server on a network. The server is not the network-visible entity. The print service will typically be a socket client on what might be called the server's request listening socket. The latter is the server's network-visible entity.

The same distinction applies to the human users of the network. They are not network-visible themselves. But a user may have an electronic mailbox on a mail server. This mailbox is network-visible and will have a network address. This distinction is quite valid, since the network does not provide any protocols for conversing with the individual user, but rather to certain applications/services through which the user can be accessed.

Entity Names

Network-visible entities may assign names to themselves, though not all NVEs have names. Such a name will be called simply an *entity name*. Entity names are character strings. A particular entity could in fact possess several names (*aliases*).

In addition to a name, an entity could also have certain attributes. For instance, a print server's request receiver might have associated with it a list of attributes of the printer, such as its type (daisy wheel, dot matrix, laser), the type and size of paper, etc. So the mapping of the entity name into its socket address might be complemented with some form of entity attributes *look-up* service.

In defining the permissible structure of an entity name, we have decided to code the attributes into the name as a separate field, known as the *entity type*.

In addition to attributes, it might prove useful to have some logically-defined location information about the entity. For instance, a given file server might belong to a particular department or building. The users of the network should be able to select a file server on the basis of its being in the vicinity or in their department, etc. Thus, the concept of a *zone* has been defined, and a *zone field* has been included in the entity name. A precise definition of the concept of a zone is provided below.

Thus, an entity name is a character string consisting of three fields: object, type, and zone concatenated together in this order with colon and @-sign separators. For example, Sidhu:Mailbox@Bandle3. Each field is an alphanumeric string of at most 32 characters. Note that by definition, all fields are *case insensitive*.

Certain meta-characters can be used in place of explicit strings. For the object and type fields, an '=' can be substituted, signifying *all possible values (wildcard)*. For the zone field, a '*' may be substituted, meaning the default value (the zone in which this node resides). If a network name does not contain any meta-characters, it is said to be fully specified. As an example, =:Mailbox@* refers to all mailboxes in the same zone as the information requester. Likewise, =:@* would mean all named entities of all types in the requester's zone. Again, Sidhu:=@* refers to all entities named Sidhu in the requester's zone regardless of their type.

Name Binding

Before a named entity can be accessed over an AppleTalk network or internet, the address of that entity must be obtained. This is done by a process known as name-binding.

Name binding may be visualized as a mapping of an entity name into its internet address, or equivalently, as a lookup of the address in a large data base, etc. (We will refer to the entity's internet address; this includes the case of a single network, where the network number field will always be equal to zero -- *unknown*).

Name binding can be done at various times. One strategy is to configure the address of the named entity into the system trying to access that entity. This so-called static binding is not appropriate for systems such as AppleTalk where the node ID can change every time a node comes up.

It is useful in this context to remember that although entities can move on a network, their names seldom change. For this reason, it is best to configure names into systems, and then use services such as NBP to bind dynamically. This may be done when the user/accessor's node is first brought up (known as early binding) or just before the access to the named entity is performed (known as late binding). Early binding runs the risk of using incorrect (out-of-date) information when the resource is actually accessed, possibly long after the user's node was brought up. However, since the binding process might be a slow one, late binding might impinge upon the response obtained by the user when accessing the named entity. Late binding is the method most appropriate when the entity is expected to *move* on the internet.

Names Directory and Names Tables

Each node maintains a *names table* (NT) containing entity name-to-entity internet address mappings (known as *NBP tuples*) of all entities in that node. Name binding is done on AppleTalk by using NBP to look up the entity's address in the *names directory* (ND). The

ND is a distributed data base of entity name-to-entity address mappings; it is the union of the individual names tables in the nodes of the internet.. The data base does not require different portions to be replicated. It can be distributed among all nodes containing named NVEs. NBP places no restriction on the use of name servers, nor are such servers necessary.

Aliases and Enumerators

NBP allows an NVE to have more than one name. Each of these aliases must be included in the NT as an independent entity.

To simplify and speed up the ability to distinguish between multiple names associated with a given socket, an enumerator value is associated with each NT entry. This is a one-byte integer, totally invisible to the clients of NBP. Each NBP implementation can develop its own scheme for generating enumerator values to be included in its NT, subject to the condition that no two entries corresponding to the same socket have the same enumerator value.

Names Information Socket

Each node implements an NBP process on a statically-assigned socket (socket number = 2) known as the names information socket (NIS). This process is responsible for maintaining the node's names table, and for accepting and servicing name lookup requests from within the node and from the network (through the NIS).

Name Binding Services

The name-binding service provides four basic services, described below.

Name Registration

Any entity can enter its name and socket number into the ND (actually into its node's NT) to make itself *visible by name*. This is done by using the name registration call to the node's NBP process.

This process must first verify that the name is not already in use. This is done by performing a name lookup in the node's zone. If the name is already in use, the registration attempt is aborted. Otherwise, the name and the corresponding socket number are inserted into the node's NT. This enters the corresponding name-to-address mapping in the ND.

When a node comes up, its names table is empty. Each network-visible entity must re-register its name(s) when restarted.

Name Deletion

A named entity should delete its name-to-address mapping from the ND when it wishes to make itself *invisible*. Reasons for doing so range from the obvious one of the entity wishing to terminate operation to sophisticated entity-specific control requirements.

The entity issues a name deletion call to the node's NBP process. The latter simply deletes the corresponding name-to-address mapping from the node's NT.

Name Lookup

Before accessing a named entity, the user (or a surrogate application) must perform a binding of the entity's name to its internet address. This is done by issuing a name lookup call to the user node's NBP process. The latter then uses NBP to perform a search through the ND for the named entity. If it is found, then the corresponding address is returned to the caller. Otherwise an *entity not found* error condition is returned.

It is possible to find more than one entity matching the name specified in the call. This is especially true when the name includes the '=' wildcard. The interface to the user must have provisions for handling this case.

Another feature of NBP is that it does not allow abbreviated names; for instance, it does not permit reference to Sidhu:Mailbox. The complete reference Sidhu:Mailbox@Bandle3 or Sidhu:Mailbox@* is required by NBP. Provisions may be made in the user interface to permit abbreviations; the interface must then *flesh out* the name before passing it on to NBP.

Name Confirmation

Name lookup performs a zone-wide ND search. More specific confirmation is needed in certain situations. For instance, if early binding was performed, the binding must be confirmed when the named entity is actually accessed. For this purpose, NBP has a name confirmation call in which the caller provides the full name and address of the entity. This call in effect performs a name search in the entity's node to confirm that the mapping is still valid.

Although a new name lookup can lead to the same result, the confirmation is more efficient in terms of total network traffic generated. It is the recommended and preferable call to use when confirming mappings obtained through early binding.

NBP on a Single Network

Name searching/look-up is quite simple on a system consisting of a single AppleTalk network.

When a user issues a name registration or lookup request to the NBP process in its node this process first examines its own names table to determine if the name is available there. If so, in the case of a registration attempt, the call is aborted with a *name already taken* result. In the case of a name lookup, the information in the names table is a partial response (there may be entities in other nodes that match the specified name).

NBP then prepares an NBP lookup packet (LkUp packet) and then calls DDP to broadcast it over the network for delivery to the Names Information Socket. Only nodes that have an NBP process will have this socket activated. In these nodes, the LkUp packet is delivered to the NBP process which searches its names table for a potential match. If no match is found, the packet is ignored. If a match is found, a LkUp-Reply packet is returned to the address from which the LkUp packet was received. This LkUp-Reply packet contains the matching name-to-address mapping (tuples) found in the replying node's names table.

The receipt of one or more replies allows the requesting NBP process to compile a list of name-to-address mappings for the original user. If the lookup was performed in response to a name registration call, then the call must be aborted since the name is already taken.

Since broadcasts are not very reliable, the requesting NBP process sends the LkUp packet several times before returning the compiled mappings, if any, to the requesting user. If no replies are received, then it is concluded that there is currently no entity using the specified name. For a name registration call, the requested name-to-address mapping is entered into the node's names table. For a name lookup call, the user is informed of a *no such entity* result.

Sending the LkUp packet several times implies that the same name-to-address mapping could be received by the requesting node several times in LkUp-Reply packets. These duplicates must be filtered out of the list of mappings. The obvious way is to compare the name strings and the address fields with each entry in the compiled list; this method, however, is inefficient. Comparison of the four-byte address fields is insufficient because of the possibility of aliases. Using the enumerator value together with the address resolves this problem, and accelerates the filtering of duplicates.

Name confirmation is similar in nature, except that the caller provides the name-to-address mapping to be confirmed. The LkUp packet is not broadcast, but is sent directly to the NIS at the specified address. This can be repeated several times to protect against lost packets or the target node being temporarily busy.

Note that on a single AppleTalk, there is only one zone. This zone should be considered to be an unnamed one (zone names originate in bridges). Any request made by a client to perform a lookup with a zone name other than '*' should be rejected with an error.

NBP on an Internet

The use of broadcast packets to perform name searching is inconvenient in internets. DDP does not allow a destination address corresponding to a broadcast to all nodes in the internet. DDP can broadcast datagrams to all nodes of any single specified network in the internet (these are said to be directed broadcasts). If NBP sent a directed broadcast to every network in the internet, the traffic generated would be considerable. For this reason, the concept of zones has been introduced.

Zones

A zone is an arbitrary subset of all the AppleTalk networks in an internet. A particular network can belong to one and only one zone. The networks in a particular zone need not be in any way contiguous or *neighbors*. The union of all zones is exactly the internet.

The concept of zones is provided to assist the establishment of departmental or other user-understandable groupings of the entities of the internet. Zones are intelligible only to the NBP (and to the related Zone Information Protocol discussed in Chapter VIII). A zone is identified by a string of at most 32 characters.

Name Lookup on an Internet

Bridges participate in the name lookup protocol of an internet. The NBP process in the requesting node prepares an NBP *broadcast request* packet (called a BrRq packet) and sends it to the NIS of A-BRIDGE. The NBP process in the bridge, in cooperation with the NBP processes in the other bridges of the internet, arranges to convert the BrRq packet into one LkUp packet for each network in the target zone of the lookup request (the exact details of this algorithm are specified in the ZIP chapter). Each of these LkUp packets is then sent

to the NIS socket as a directed broadcast on the corresponding network. The replies are returned to the original requester as before.

The important point is that non-bridge nodes do not need to know anything about zones. The process of generating a zone-wide broadcast is done by the collection of bridges. The latter must of course jointly have a complete mapping of zone names into the list of corresponding networks. The establishment and maintenance of these mappings is the purpose of the Zone Information Protocol discussed in Chapter VIII.

Note that on an internet, nodes performing lookups in their own zone will receive LkUp packets from themselves (via a bridge). The node's NBP process must not respond to these!

Note that the name registration process of NBP tries to ensure that the same name is not used by more than one entity in a given zone. However, if another network is *moved* into the zone by changing its zone name, then it is possible to end up with two or more entities (in the new zone) with the same name.

NBP Interface

Four calls provide to the user all the functionality of name-binding; they are described below.

Note that all calls to NBP take as a parameter an entity name. It is only the lookup call, however, where the zone name is meaningful. In all other calls, it is required, for consistency, that the zone name be set to '*'.

Registering a Name

This call is used by an NBP client to register an entity name and its associated socket address. Meta-characters are not allowed in the object and type fields; the zone name field should be equal to '*'.

- Call Parameters:

- entity name
 - socket number

- Returned Parameters:

- result code: success
 - failure (name conflict, invalid name or socket)

Although this is not a required feature of NBP it is advisable for the implementation of this call to verify that the socket is in fact open.

Removing a Name

This call removes an entity name from the node's names table. Meta-characters are not allowed in the object and type fields; the zone name field should be equal to '*'.

- Call Parameters:

entity name

- Returned Parameters:

result code: success
failure (name not found)

Name Lookup

This call performs the mapping between entity name and address. Meta-characters are allowed in the name to make the search as general as needed. It is possible for more than one address to match the call's entity name. For each match, this call returns the name and its internet address. The names contain fully specified object and type fields; the zone name, however, will always be '*', regardless of zone specified.

- Call Parameters:

entity name
maxMatches

- Returned Parameters:

result code: success
failure (name not found)
list of entity names and their corresponding internet addresses

The parameter maxMatches is a positive integer that specifies the maximum number of matching name-to-address mappings needed. This is useful if wildcards are used by the caller in the entity name parameter.

Confirming a Name

This call confirms a caller-supplied mapping between entity name and address. Meta-characters are not allowed in the object and type fields; the zone name field should be equal to '*'.

- Call Parameters:

entity name,
socket address

- Returned Parameters:

result code: success (mapping still valid)
wrong-socket (net and node number valid, socket number invalid)
failure (mapping invalid)

new socket (only if wrong-socket error code)

NBP Packet Formats

NBP packets are identified by a DDP protocol type field of 2. There are three different types of NBP packets: BrRq, LkUp and LkUp-Reply. The format is as shown in Figure VI-1. It consists of an NBP header followed by the name-address pairs (known as NBP tuples); the various fields are described below.

Control

The high-order four bits of the first byte of the header are used to indicate the type of NBP packet. The values are 1 for BrRq, 2 for LkUp and 3 for LkUp-Reply.

Tuple Count

The low-order four bits of the first byte contain a count of the number of NBP tuples in that packet. BrRq and LkUp packets carry exactly one tuple (the name being looked up or confirmed). The tuple count field for these packets is always equal to 1.

NBP ID

In order to allow for multiple pending lookup requests from a given node, an 8-bit ID is generated by the NBP process issuing the BrRq or LkUp packets. The LkUp-Reply packets must contain the same NBP ID as the LkUp or BrRq packet to which they correspond.

NBP Tuple

The format of the NBP tuples, the name-address pairs, is shown in Figure VI-2. The tuple consists of an entity's name, its 4-byte internet address, and a one-byte enumerator field. The address field appears first in the tuple. The fifth byte in a tuple is the enumerator field, followed by the entity name. This consists of three string fields: one each for the object, type, and zone names. Each of these strings consists of a leading 1-byte string length followed by up to 32 string bytes. The string length represents the number of bytes/characters in the string. The three strings are concatenated without any intervening padding.

The enumerator field is included to handle the situation where more than one name has been registered on the same socket. It should be noted that NBP specifically permits this use of aliases (or alternately, the use of a single socket by more than one NVE). In this case, each alias is given a unique enumerator value, kept in the NT along with the name-address mapping. The enumerator field is not significant in a LkUp or BrRq packet, and is ignored by the recipient of these packets.

In BrRq and LkUp packets (which carry only a single tuple) the address field contains the internet address of the requester; this allows the responder to properly address the LkUp-Reply datagram. In a LkUp-Reply packet, the correct enumerator value must be included in each tuple. This value is used for duplicate filtering.

Since normal nodes (and hence their NBP processes) are usually unaware of zone names, specifically their own zone name, tuples in LkUp-Reply packets do not specify a zone name. The zone names in these tuples are always equal to '*', regardless of the zone in

which the lookup is performed. The requestor will know the zone name of these responses, since it must be the zone he asked for in the lookup request.

Note that in a LkUp or BrRq request, a null zone name (length byte equal zero) is equivalent to '*'.




Figure VI-1 NBP Packet Format




Figure VI-2 NBP Tuple

VII. AppleTalk Transaction Protocol (ATP)

About AppleTalk Transaction Protocol

The fundamental purpose of reliable transport protocols is to provide loss-free delivery of client packets from source socket to destination socket. Various features can be added to this basic service to obtain characteristics appropriate for specific needs. Transport characteristics can be added by clients or built into standard value-added transport services.

The AppleTalk Transaction Protocol (ATP) is a transport protocol that satisfies the transport needs of a large variety of peripheral devices, and the transaction needs for more general networking on AppleTalk. Particular attention has been paid to ease of implementation, so nodes with tight memory space restrictions will be able to support a sufficient subset of ATP.

Transactions

Often a socket client must request the client of another socket to perform a particular higher level function and to report the outcome. This interaction, between a requester and a responder, is termed a *transaction*.

The basic structure of a transaction, in the context of a network, is illustrated in Figure VII-1. The requester initiates the transaction by sending a transaction request, TReq, from the requester's socket to the responder's socket. The responder executes the request and returns a transaction response, TResp, reporting the transaction's outcome.

At-Least-Once Transactions (ALO)

This basic process must be performed in the face of various situations inherent in the loosely-coupled nature of networks, for example:

- TReq is lost in the network
- TResp is lost or delayed in transit
- The responder *dies* or becomes unreachable from the requester

There could be several transaction requests outstanding, and the requester must be able to distinguish between the respective responses received by it over the network. This can be done by using a transaction identifier (TID), sent with the request. The response must contain the same TID as the corresponding request. The TID, in a sense, *binds* together the request and the response portions of a transaction.

In any of the above three error situations, the requester will not receive the transaction's response and must conclude that the transaction was not completed. A recovery procedure must then be activated at the requester. This consists of a timeout and an automatic retry mechanism. If the timer expires in the requester and the response has not been received, the requester retransmits the TReq (see Figure VII-1). This process is repeated until a response is received or until a maximum retry count is reached. If the retry count hits its maximum value then it is concluded that the responder is either *dead* or unreachable, and the transaction requester (the ATP client at the requester) is so notified.

This mechanism does its best to ensure that the transaction request is executed at least once. Such a mechanism is adequate if the request is essentially idempotent, that is, repeated execution of the request is the same as executing it once. (An example of an idempotent transaction is asking a destination station to identify itself.)

Exactly-Once Transactions (XO)

Nevertheless, with this recovery mechanism, lost or delayed responses could cause the transaction to be executed more than once. If the request is not idempotent, serious damage could result. For such requests it is desirable to have a transaction service which ensures transaction execution once and exactly once.

Whether the at-least-once or the exactly-once level of service is appropriate can only be determined by the transaction requester.

The basic technique for implementing an exactly-once transaction protocol is as follows (see Figure VII-2). The responder maintains a transactions list of all recently received transactions. Upon receiving a TReq the responder searches through this list to determine if the request has already been received (duplicate transaction-request filtering). A newly received request is inserted into the list and is then executed. After it has been executed, the response is sent back and a copy of the response is attached to the transaction's entry in the transactions list. Upon receiving a duplicate request for which a response has already been sent, the responder retransmits the response. If a duplicate request is received, and a response has not been sent out yet (the request is still being executed), then ATP ignores the duplicate request.

Upon receiving a TResp, the requester should return a transaction release (TRel) packet to release the request from the responder's transactions list. If this TRel gets lost, then the request would stay in the list *forever*. To eliminate this situation, the responder *timestamps* a request before inserting it in its list. The list is checked periodically and requests that have been in it for too long are eliminated.

This method of filtering duplicate requests by consulting a list of recently received transactions is quite effective in ensuring exactly-once service in most environments. However, it does not guarantee exactly-once service in *all* environments. In fact, if the situation is such that packets are guaranteed, if they arrive at all, to arrive in the order in which they were sent (on a single AppleTalk network), then this technique is fully effective.

However, in an internet environment, due to multiple paths from source to destination and different transmission delays on these paths, packets may arrive at their destination in an order different from the one in which they were sent. Thus, unusual situations such as the one illustrated in Figure VII-3 can arise. Here, the original transaction response was delayed long enough in the internet to provoke a retransmission of the request.

Furthermore, the TRel sent by the requester upon receiving the response arrives before the retransmitted request. The responder releases the request from the list upon receiving the TRel, and thus the retransmitted request cannot be filtered out as a duplicate. In fact, with a connectionless protocol it is difficult to guarantee exactly-once service under internet conditions. ATP XO mode does, however, greatly reduce the amount of work a connection-based client must perform to realize true exactly-once service.

ATP XO guarantees that, if a duplicate request is received by the responder, the transaction has already been completed and the request can be ignored. Clients desiring absolutely

guaranteed exactly-once service over an internet should employ some form of simple sequence number checking in addition to ATP to achieve guaranteed internet exactly-once service (an example of this is detailed in the chapter on the Printer Access Protocol).

Note that ATP XO should be considered an optional part of ATP. Nodes that do not require exactly-once service do not need to implement it. It should be kept in mind that higher level protocols, may require ATP XO (for example, the Printer Access Protocol and the AppleTalk Session Protocol).

Multi-Packet Responses

ATP uses the basic idea that a transaction is a request issued by a client in a requesting node to a client in a responding node. The client in the responding node is expected to service the request and generate a response. It is assumed that these clients have some method of unambiguously identifying the data/operation sought in the request (e.g. read a disk block, block number, etc.).

This is a very simple model, in principle sufficient for all interactions. The difficulty is that the underlying network places a size restriction on the packets that can be exchanged. Thus, for instance, the transaction response might not fit in a single packet. For this reason we look upon the transaction request and response as *messages* (not packets). Although ATP restricts its transaction requests to single packets, it allows the transaction response message to be made up of several packets, which, of course, bear a sequential relationship to one another. When the requesting node receives all the response packets (the complete response message), the transaction is considered complete and the response is delivered as a single logical entity to the ATP client (the transaction requester).

ATP supports both *at-least-once* and *exactly-once* modes as client-elected options.

If the 'at-least-once' case is chosen then ATP handles timeouts and retransmission of requests but does not handle retransmission of responses. In this case, if request actions are not idempotent, it is up to the responding client to handle retransmission of responses to duplicate requests.

The maximum size (number of packets) of a transaction response message is limited to 8 packets. The maximum amount of data in an ATP packet (request or response) is 578 bytes. This limit is derived from the DDP maximum packet size of 586 bytes minus ATP's header size of 8 bytes.

Transaction IDs

Transaction IDs are generated by the requester and sent along with the TReq packet. An important design issue is the size (16 bits for ATP) of these IDs. This is a function of the rapidity with which transactions are generated and of the maximum packet lifetime (MPL) for the complete network system. The longer the MPL, the larger the TID must be. Similarly, if transactions are generated rapidly, then the TIDs must again be larger. The basic problem is that the TID being of finite size wraps around and there is the danger that for a particular value an old packet stored in some internet router may arrive later on and be accepted as a valid packet.

For a single AppleTalk network, the time taken for exchanging a TReq and a TResp is bounded (by the ALAP) to be greater than about 2.5 msec. Thus there can be no more than

400 transactions per second. From this point of view a single byte TID would allow half a second or so for wraparound of the TID.

However, with network interconnection through store-and-forward internet routers, the impact of MPL (of the order of 30 seconds) makes a one-byte TID inadequate. A 16-bit TID would increase the wrap around time to be approximately two minutes. This obviates concerns about old retransmitted transaction requests and responses "sneaking-in" due to wraparound.

Later in this chapter, the issue of generating transaction identifiers will be revisited to account for another subtle but important characteristic of ATP, namely transactions with infinite retries.

ATP Bitmap/Sequence Number

Every ATP packet includes in its header a bitmap/sequence number. This field is 8 bits wide. ATP handles lost or out-of-sequence response packets by using this bitmap. The significance of this field depends on the type of ATP packet (TReq, TResp or TRel).

In TReq packets this field is known as the transaction bitmap. The basic idea behind the use of this field is that the requester reserves enough buffers for the expected transaction response, and then sends out the TReq packet indicating to the responder the number of buffers reserved. This is done by setting a bit in the TReq packet's bitmap, for each reserved buffer. The responder can then examine the TReq packet's bitmap and determine the number of packets the requester is expecting to receive in the transaction response message.

In TResp packets this field is known as the *ATP sequence number*. The value of this field in the TResp packet is an integer (in the range 0 to 7), indicating the sequential position of that packet in the transaction response message. The requester can use this value to put the received response packet in the appropriate response buffer (even if the response packet is received *out-of-sequence*) for delivery to the transaction requester (ATP client). Furthermore, the requester clears the corresponding bit in its copy of the transaction bitmap.

The actual transaction response message may turn out to be smaller than was expected by the requester. Thus a provision is made in the response packet's header to signal *end-of-message* in the last response packet when it is sent out by the responder. Upon receiving a response packet with the end-of-message indication, the requester must clear all bits in its copy of the transaction bitmap corresponding to higher sequential positions. Note that this end-of-message signal is internal to ATP; the responding client tells ATP to set it, but it is not necessarily delivered to the requesting client and should not be used for higher-level communications (e.g. as an end-of-file indicator).

If the requester's retry timeout expires and the complete transaction response has not been received as yet (indicated by one or more bits still set to one in the requester's transaction bitmap), then a TReq is sent out again with the current value of the transaction bitmap and the same TID. Thus only the missing transaction response packets are requested again.

This mechanism is illustrated in Figure VII-5 where a requester issues a TReq indicating that it has reserved six buffers for the response. For instance, the request might be for six blocks of information from a disk device. The TReq packet would have in its ATP data part the pertinent information: what file, which six blocks, etc. ATP builds the request

packet and sets the least significant six bits in the bitmap. When the responder receives this request packet, it examines the request's ATP data and its bitmap and thus determines the type and range of the request to be serviced. The six blocks are fetched from disk, passed to the ATP layer in the device node and sent back to the requesting node, each in a separate packet with its sequence number indicating the position in the response.

The example illustrated assumes that the third response packet is lost in the network. Thus the retry timeout will expire in the requester, which then retransmits the original request (transparently to the ATP requesting client) but with a bitmap reflecting only the missing third response packet.

Note that single packet request-response transactions are simply the degenerate case in which the transaction request has only one bit set in its bitmap. If two nodes wish to communicate in this manner, very little extra packet overhead is added by the protocol.

Responders with Limited Buffer Space

A potential difficulty with exactly-once transactions is that a responder might not have enough buffer space to hold the entire transaction response message until the end of the transaction (for example, receipt of a TRel).

ATP provides a mechanism for such responders to reuse their buffers, through a confirmation of response packet delivery. This is done by piggy-backing, in a response packet, a request to Send Transaction Status (STS). The requester, upon receiving such a response packet, immediately sends out a TReq with the current bitmap (indicating which response packets have still not been received and hence providing a way to determine which have already been received). The responder can then use this bitmap to free buffers holding already-delivered response packets.

Two user-interface issues arise in connection with the STS bit. The retransmitted TReq would normally be filtered by ATP XO as a duplicate, so ATP must provide some way of delivering the updated bitmap to the user without delivering the duplicate request. Related to this is the fact that, in an internet, TReq's can get out of order; if a duplicate TReq is ever received whose bitmap indicates that fewer responses have been received than a previous TReq indicated, it should be ignored as a delayed duplicate and should not be passed to the user.

Figure VII-6 illustrates the use of STS with an example in which the responder, with two buffers, services a request for a seven packet response.

TReq packets sent in response to an STS do not consume the retry count, but do reset the Retry Timeout.

ATP Packet Format

The format of an ATP packet is illustrated in Figure VII-7. It consists of an 8-byte ATP header plus up to 578 ATP data bytes.

The first byte of the ATP header is used for command/control information (CCI). The two high-order bits of the CCI are the packet's function code. These two bits are encoded as follows:

- 01 -- TReq packet
- 10 -- TResp packet
- 11 -- TRel packet.

The EOM bit is set in a TResp packet to signal that it is the last packet in the transaction's response message. The XO bit must be set in all TReq packets that pertain to the exactly-once mode of operation of the protocol. The STS bit is set in TResp packets to force the requester to retransmit TReq immediately. The remaining three bits of CCI must always be zero.

The eight bits immediately following the command/control field contain the ATP bitmap/sequence number. The packets comprising the transaction response message are assigned sequence numbers 0 through 7. The sequence number (encoded as an integer) is sent in the ATP bitmap/sequence number field of the corresponding response packet.

In the case of a transaction request packet, a bit of the bitmap is set to one for each expected response packet. The least significant bit corresponds to the response packet with sequence number 0, up through the most significant bit which corresponds to sequence number 7.

The third and fourth bytes of the ATP header contain the 16-bit transaction ID. TIDs are generated in the ATP requester, and are incremented from transaction to transaction as unsigned 16 bit integers (the value zero is permitted).

The last four bytes of the ATP header are not examined by ATP, but contain user data. As such, they should, strictly speaking, not be considered part of the ATP header. However, they can be used by an ATP client to build a simple header for a higher level protocol. The reason they have been separated out is to allow an implementation of ATP that handles the complete ATP response message's data in an assembled, contiguous, form, without interposed higher level headers. ATP-Client interfaces must build appropriate mechanisms for exchanging these four *user bytes* independently of the data. It should also be noted that the ATP user bytes contained in a TRel packet are not significant and clients should not use them in any way.

ATP Interface

The interface to ATP is made up of the five calls described below. It is convenient to visualize the ATP package (an implementation of ATP) as consisting of two parts: the ATP Requester and the ATP Responder. The calls are discussed in the context of each end.

It is not appropriate in this specification of the protocol to detail the interface, as several aspects are implementation dependent. The description below is in general terms, adequate for establishing the characteristics of the ATP service to the next higher layer.

Remember that we neither assume nor require the availability of a multiprocessing environment in the network node. Our descriptions of the various interface calls have been written in a generic form indicating parameters passed by the caller to the ATP implementation, and results and outcome codes returned by the latter. The result codes and their interpretation depend on the specifics of the implementation of a call. If the call is issued synchronously, the caller is blocked until the call's operation has been completed or aborted, then the returned parameters become available when the caller is unblocked. In the case of asynchronous calls, a call completion mechanism is activated when the operation completes or aborts. Then the returned parameters become available through the completion routine mechanism.

We envision at least two kinds of interfaces: a packet-by-packet passing of response buffers to and from ATP, and a response message (that is, in a contiguous buffer) interface. These are analogous to the familiar packet stream and byte stream interfaces available for data stream protocols. However, implementors are completely free to provide any type of interface they consider appropriate.

Send a Request

The transaction requester (ATP client) issues this call to send a transaction request. It must supply several parameters with the call. These include the address of the destination socket, the ATP data part and user bytes of the request packet, buffer space for the expected response packets, and whether the exactly-once mode of service is required or not. In addition the caller specifies the duration of the retry timeout to be used and the maximum number of retries. There must be a provision in the interface for the caller to indicate *infinite retries*, that is, keep trying the request until a response is obtained. It may also be desirable for the caller to be able to optionally specify the socket through which the request should be sent (or ATP could pick one for him).

• Call Parameters:

- transaction mode (exactly-once or at-least-once)
- transaction responder's address (network number, node ID and socket number)
- ATP request packet's data part and its length
- ATP user bytes
- expected number of response packets
- buffer space for the transaction response message
- retry timeout in seconds
- maximum number of retries
- [optional] socket through which to send the request

• Returned Parameters:

- result code: (success, failure)
- number of response packets received
- user bytes from responses

An outcome code of failure is returned if ATP has exhausted all retries and a complete response has not been received. Note that no error is returned if the caller requested exactly-once service but the responder does not support it -- the request will be executed at least once.

An outcome code of success is returned whenever a complete response message has been received. A complete response is said to have been received if either of the following occurs: (1) all response packets originally requested have been received, or (2) all response packets with sequence number 0 to some integer *n* have been received and packet *n* had the EOM bit set.

In either case, the actual number of response packets received is always returned to the requesting client. A count of zero should indicate that the other end did not respond at all. In the case of a nonzero count, the client can examine the response buffers to determine which portions of the response message were actually received and to detect missing pieces for higher level recovery, if so desired.

Open a Responding Socket

An ATP client uses this call to instruct ATP to open a socket (well-known or dynamically assigned) for the purpose of receiving transaction requests. If well-known, the client passes the socket number to ATP; otherwise the dynamically assigned socket number is returned to the client.

When opening this socket, the client is in effect opening a *transaction listening socket*. The call allows the socket to be set up so that requests are accepted from a specified network address (provided in the call). This address can include a zero in the network number, node ID, or socket number fields to indicate that any value is acceptable for that field.

- Call Parameters:

- transaction listening socket number (if well-known)
 - admissible transaction requester address (network number, node ID, and socket number)

- Returned Parameters:

- result code (success, failure)
 - local socket number (if dynamically assigned)

Note that this call does not set up any buffers for the reception of transaction requests. That is done by issuing the Receive-a-Request call.

Close a Responding Socket

This call is used to close a socket previously opened with the Open-a-Responding-Socket call.

- Call Parameters:

- transaction listening socket number

- Returned Parameters:

- result code (success, failure)

Receive a Request

The transaction responder issues this call to set up the mechanism for actual reception of a transaction request through an already-opened transaction responding socket.

- Call Parameters:

- local socket number on which to listen
 - buffer for receiving the request

- Returned Parameters:

- the received request's ATP data
 - received request's user bytes
 - transaction ID

transaction Requester's Address (network number, node ID and socket number)
bitmap
XO indication

Send a Response

When a transaction responder has finished servicing the request, it issues a Send-a-Response call to send out one or more response packets. ATP will send out each response buffer with the indicated transaction ID and a sequence number indicating the position of the particular response packet in the response message. There are many different ways of implementing this call; our description is generic.

- Call Parameters:

local socket number (the responding socket)
transaction ID
destination internet address (Network number, node ID and socket number)
transaction response message/packets (ATP data part)
transaction user bytes (up to eight sets)
descriptors to determine the sequence numbers of the response packets
EOM and STS control

- Returned Parameters:

result code: (success, failure)

ATP State Model

The following description provides a sufficiently precise statement of the protocol internals for the purpose of implementing the protocol. It is not a formal specification, but an aid for protocol implementers.

The description is presented in terms of the actions to be taken in response to all possible events. Certain special terms are employed in the description, and we start with a discussion of some of these.

First of all, the ATP requester must maintain all information necessary for retransmitting an ATP request and for receiving its responses. This is referred to by us as the Transaction Control Block (TCB). More specifically, this would contain all the information provided by the transaction requester in the Send-a-Request call, plus the TID, the request's bitmap and a response-packets-received counter. With each request and thus with each TCB we associate a retry timer, used to retransmit the request packet to recover from loss of request or response packet situations.

In the second place, the ATP responder must maintain a Request Control Block (RqCB) for each Receive-a-Request call issued by a client in that node. This block contains the information provided by that call including all pertinent data as to the buffers and delivery-to-client mechanism (implementation dependent).

A third data structure, the Response Control Block (RspCB) is needed only in nodes implementing the exactly-once mode of operation. It holds the information needed to filter duplicate requests and to retransmit response packets in response to these duplicates. We associate a release timer with each RspCB. This timer is used to release the RspCB in the

event that the release packet sent by the requester is lost. The transactions list mentioned earlier consists of these RspCBs.

The release timer is specified to be 30 seconds long. Note that the release timer is started as soon as a RspCB is set up (i.e. upon receiving the TReq). It is reset every time a corresponding TResp is sent out. This implies that the responding client must send out the first TResp within a thirty second interval of the TReq's arrival, and subsequent TResps at a maximum separation of thirty seconds from each other. Failure to do so will result in the RspCB being destroyed and a duplicate request being delivered to the responding client.

ATP Requester

Send-a-Request call issued by a transaction requester in the node:

- A. Validate the following call parameters:
 - number of response packets should be at most 8
 - the ATP request's data should be at most 578 bytes long;if either parameter is invalid, then reject the call;
- B. Create a TCB:
 - insert the call parameters in it
 - clear the response-packets-received counter
 - insert the Retry Count into the TCB;
- C. Generate a TID:
 - this TID must be such that the packets of the new transaction will be correctly distinguished from those of other transactions; details of TID generation are discussed at the end of this chapter;
 - save the TID in the TCB;
- D. Generate the bitmap for the request packet and save it in the TCB;
- E. Prepare the ATP header:
 - insert the TID and the bitmap
 - set the function code bits to binary 01
 - {only for systems implementing the *exactly once mode* of operation}:
if the caller requested *exactly once mode* set the XO bit;
- F. Call DDP to send the ATP request packet (disregard any error returned);
- G. Start the request's retry timer.

Retry timer expires:

- A. If Retry Count = 0 then:
 - set outcome code to *failure*
 - notify transaction requester (the client in the node) of the outcome
 - destroy the TCB;

- B. If Retry Count $\neq 0$ then:
- decrement the retry count (if not *infinite*)
 - change the bitmap in the ATP request's header to current value in the TCB
 - call DDP to retransmit the request packet (ignore errors)
 - start the retry timer.

ATP Response Packet received from the network (i.e. from DDP):

- A. Use the packet's TID and source address to search for the TCB
- B. If a matching TCB is not found then ignore the packet and exit
- C. If a matching TCB is found then check the packet's sequence number against the TCB's bitmap to determine if this response packet is expected. The packet is expected if the bit corresponding to the response packet's sequence number is set in the TCB's bitmap. If the packet is not expected then ignore it and exit
- D. If the response is expected then:
- Clear the corresponding bit in the TCB bitmap
 - Set up response packet's ATP data for delivery to transaction requester
 - Increment the response packets counter in the TCB
- E. If the packet's EOM bit is set then clear all higher bits in the TCB bitmap
- F. If the packet's STS bit is set then:
- Call DDP to re-send the request with the current TCB information
 - Reset the retry timer for the request;
- G. If the TCB bitmap = 0 then: {a complete response has been received}
- Cancel the retry counter
 - Set the outcome code to *success*
 {only if exactly-once mode is implemented}
 if the transaction is of exactly-once mode (determined by
 examining the TCB) then call DDP to send a TRel packet
 to the responder
 - Notify the transaction requester
 - Destroy the TCB

ATP Responder

Open-a-Responding-Socket call issued by a Transaction Responder in the Node:

- A. If caller specifies a well-known socket then call DDP to open that socket
 else call DDP to open a dynamically assigned socket;
- B. If DDP returns with error then set result code to the error;

C. If DDP returns without error then

- set result code to *success*
- save socket number and the admissible transaction requester address in an ATP responding sockets table;

D. Return to caller .

Close-a-Responding-Socket call issued by a transaction responder in the Node:

A. Call DDP to close the socket

B. Release all RqCBs, and, for systems supporting exactly-once mode of operation, release all RspCBs (and cancel all release timers), if any, associated with the socket

C. Delete the socket from the ATP responding sockets table

Receive-a-Request Call issued by the transaction responder in the node:

A. If the specified local socket is not active then return to caller with error

B. Create a RqCB and attach it to the socket

C. Save the call's parameters in the RqCB

Send-a-Response call issued by transaction responder in the node:

A. If the local socket is invalid OR response data lengths are invalid then return to caller with error

B. {only if *exactly once mode* is implemented} Search for a RspCB matching the call's local socket number, TID, and transaction requester address (destination of the ATP Response). If found then save the response attached to the RspCB (for potential retransmission in response to duplicate requests received subsequently), and restart release timer

C. Send the response packets through DDP, setting the ATP header of each with function code binary 10, with the caller supplied TID, with the correct sequence number for the packet's sequential position in the response message, with the EOM flag set in the last response packet (and the STS flag if requested). Ignore any error returned by DDP

*Release timer expires: {only if *exactly once mode* of operation is implemented}*

a. destroy the RspCB and release all associated data structures.

ATP Request Packet (TReq) received from DDP:

- A. {only if *exactly once mode* is implemented} If the packet has its XO bit set and a matching RspCB (that is the packet's source and destination addresses are the same as those saved in the RspCB and the packet's TID is equal to the one saved in the RspCB) exists then:
- retransmit all response packets
 - restart the release timer
 - return the bitmap to the client if a previous response had the STS bit set
 - exit
- B. If a RqCB does not exist for the local socket or if the packet's source address does not match the admissible requester address in the RqCB then ignore the packet and exit
- C. {only if *exactly once mode* is implemented} If packet's XO bit is set then create a RspCB (save therein the request's source and destination addresses and its transaction ID) and start its release timer
- D. Notify the client about the arrival of the request and destroy the corresponding RqCB

ATP Release Packet (TRel) received from DDP: {only if exactly-once mode is implemented}

- A. Search for a RspCB that matches the packet's TID and source and destination addresses. If not found then ignore the release packet;
- b. If a matching RspCB is found then:
- Destroy the RspCB and all release associated data structures
 - Cancel the RspCB's release timer.

Some Optional ATP Interface Calls

In certain instances the clients of ATP might use certain contextual information to enhance their use of ATP through some additional interface calls. Here are two examples. These calls have been found useful in implementing certain higher level protocols, but are completely optional for a given ATP implementation.

Release a RspCB

In the first case two clients of ATP are communicating with each other, using the exactly once mode, and have decided to have at most one outstanding transaction at a time. Client A calls ATP to send a TReq packet to client B. B sends back the response. A, upon receiving the response, sends out a second request (but no release packet). The second request packet upon being received by B signals that the response to the previous request has been received by A. Now B could simply call its ATP responder and ask it to release the previous transaction's response control block. This needs a Release-RspCB call to the ATP responder. The parameters of this call are fairly obvious in this instance.

Release a TCB

Another instance arises when a client A wishes to send data to client B. A must first inform B of this intention, and thus allow B to request the data. To this end, A can send a TReq to B to signal "I want to write N bytes of data to you, please ask me for it on my socket number S". Instead of sending a TResp to this packet, B could just send a TReq to A's socket S asking for the data. The reception by A, on socket S, of B's request implies that A's original request has been received by B. A could call its ATP requester and ask it to eliminate the previous transaction's TCB. This needs a Release-TCB call to the ATP requester. This call could also be used by the requesting-end client to cancel an outstanding ATP request at any time -- for instance to abort an infinite retry request.

Details of these calls are implementation-dependent and do not affect the ATP protocol per se. We mention these as interesting variants that might be implemented by certain clients who wish to reduce the traffic generated on the network.

Wraparound and Generation of Transaction Identifiers

Earlier in this chapter it was noted that transaction IDs being of finite size wrap around and that there is the danger that for a particular value an old packet stored in some internet router may arrive later and be accepted as a valid packet in a later transaction using the same identifier. On the basis of a maximum packet lifetime estimate of 30 seconds, it was determined that this problem could be avoided if the TID were 16 bits long (wraparound is estimated to take approximately two minutes).

There is yet another important problem of this nature that stems from the fact that a particular transaction might take much more than two minutes to complete. For instance, if the request is to search through an *encyclopedia* for all references to a particular piece of information, this might take a very long time indeed (several minutes). In this case, the transaction ID could wrap around and another transaction is then issued with the same TID, leading to the same dilemma as the one mentioned in the previous paragraph. ATP does not prohibit operations of this sort. In fact, it is because such transactions might be made, that the specification of the length of the ATP retry timer and maximum retry count is left up to ATP's transaction requesting client.

In the same vein, ATP allows the requester to issue an ATP transaction with maximum retry count set to *infinite*. In this case, ATP continues to retransmit the transaction request until a reply is received. If a reply is not sent, then the transaction will be retransmitted infinitely often. This leads to the same TID wraparound problem. In fact, some protocols like the AppleTalk Session Protocol use infinite-retry transactions for the purpose of tickling the other end of a session (that is, informing it that this end is alive). These "tickle" transactions never receive a reply.

A properly implemented ATP will function correctly in the face of these wraparound scenarios. There are two aspects to this: the use of TIDs to distinguish between transactions and the generation of TIDs.

The ATP requesting end generates the TID when it is asked by a client to send a transaction request. At that time, it creates a TCB and saves several pieces of information in it. These include the number of the local socket through which the transaction is being sent, the complete internet address of the responding socket to which it is being sent, and the just generated TID. This information is saved in order to ensure a correct match of the response packets with the transaction. When the ATP requesting end receives a transaction response

it identifies the corresponding request by looking for a TCB whose saved information matches the response packet's TID and the packet's source and destination socket addresses.

Thus TID wraparound by itself is not dangerous except if it causes the simultaneous existence of two or more transactions (that is, TCBs) with the same TID *and* the same requesting and responding socket addresses. This observation allows the specification of an algorithm for generating TIDs that ensures that wraparound has no ill effects. This is as follows:

```
{ Algorithm used by ATP Requesting end to generate TID for a new transaction }
new_TID := last_used_TID;
Not_In_Use := TRUE;
REPEAT
    new_TID := (new_TID + 1) modulo 216;
    Search all TCBs on the local requesting socket, and if any one of these has
        (new_TID = TCB's TID) then set Not_In_Use := FALSE;
UNTIL Not_In_Use;
{ At this point new_TID has the newly generated transaction identifier }
last_used_TID := new_TID;
```

Note that this algorithm ignores the TCB's destination socket address (that is, does not compare it with the address of the new transaction's destination). This simplification of the algorithm does not however in any way reduce its effectiveness in preventing the wraparound problem.


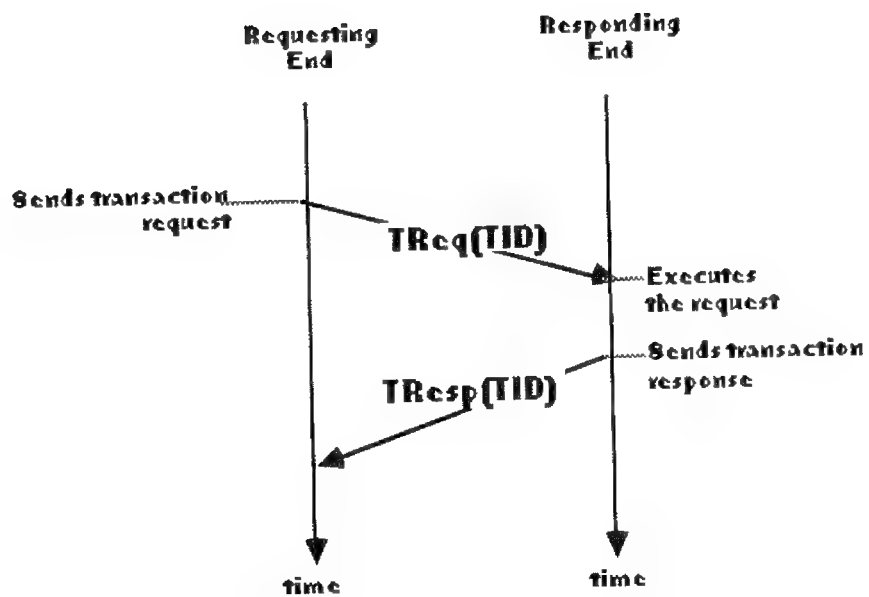


Figure VII-1: Transaction

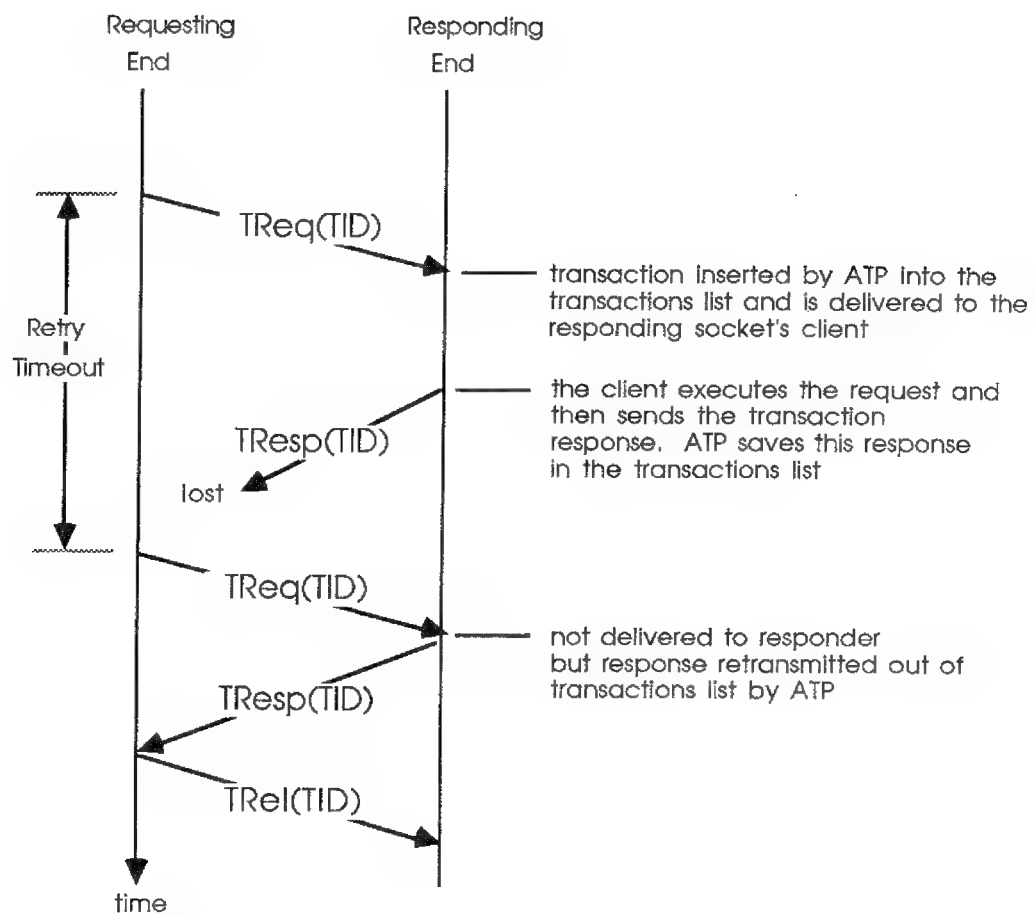


Figure VII-2: Exactly-Once Transactions

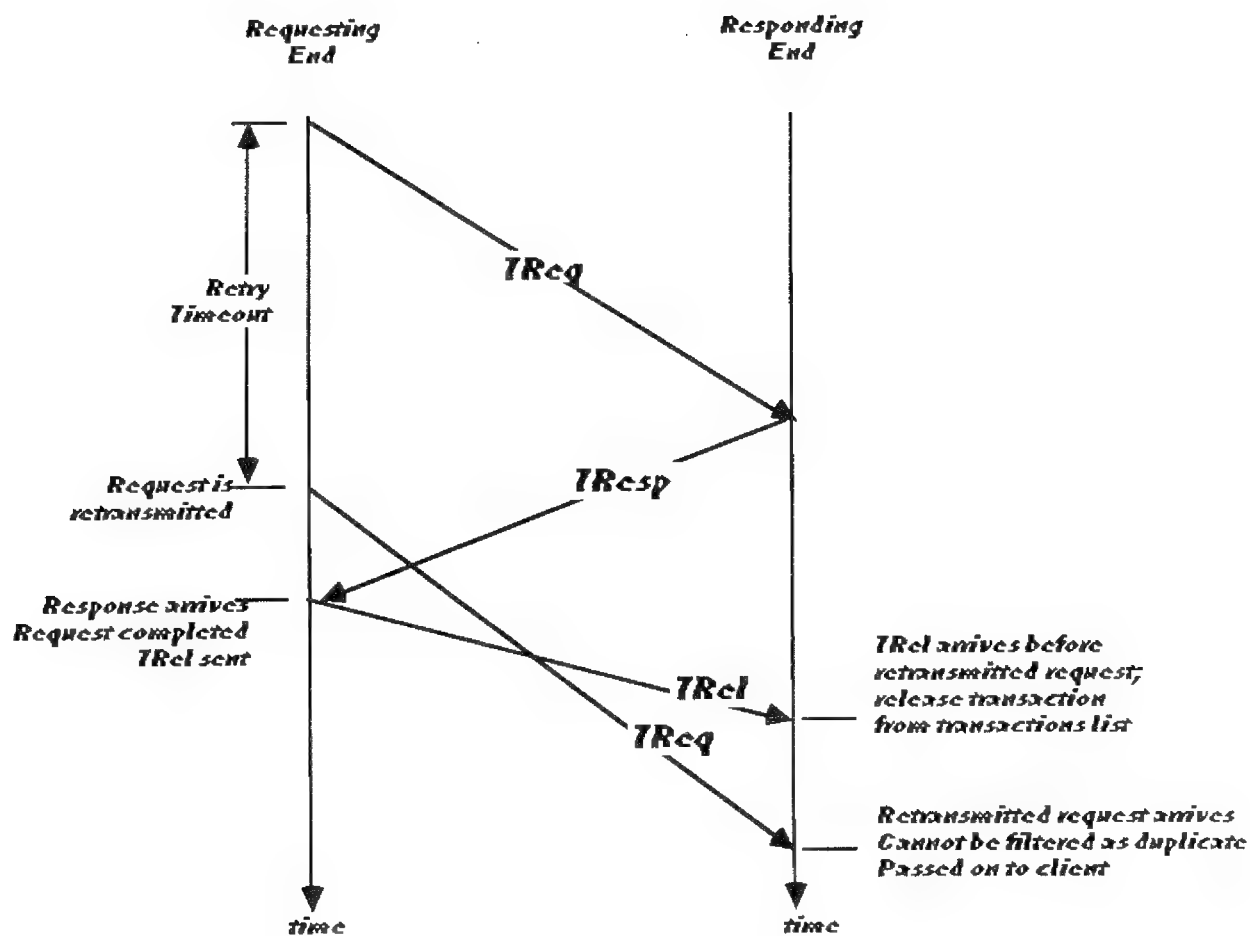


Figure VII-3: Duplicate delivery of Request in

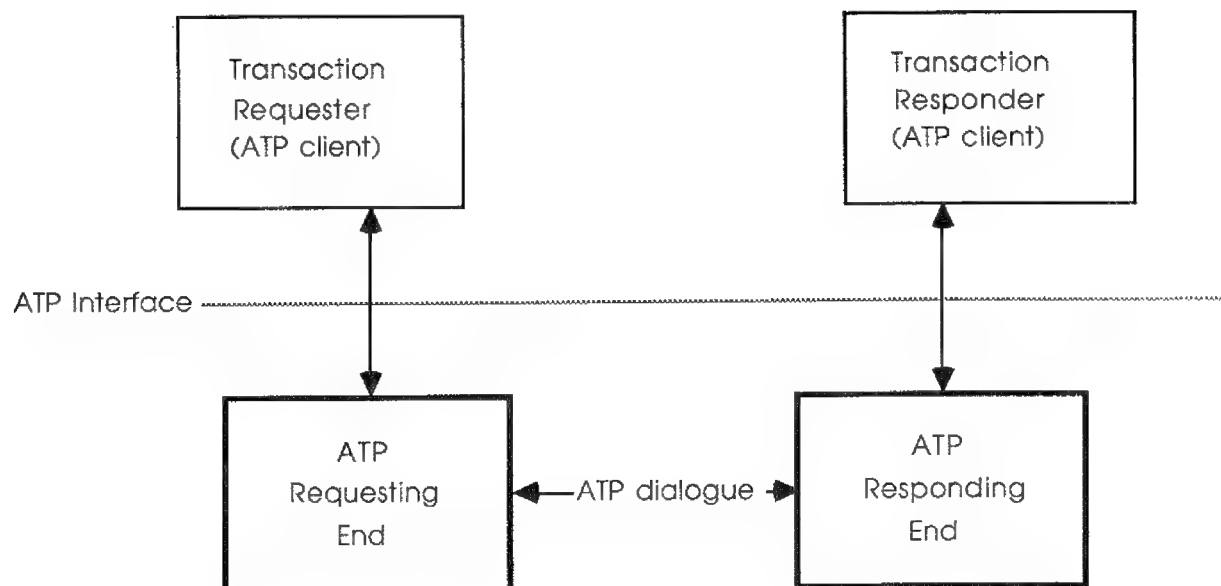


Figure VII-4: ATP Terminology




Figure VII-5: Multiple-Packet

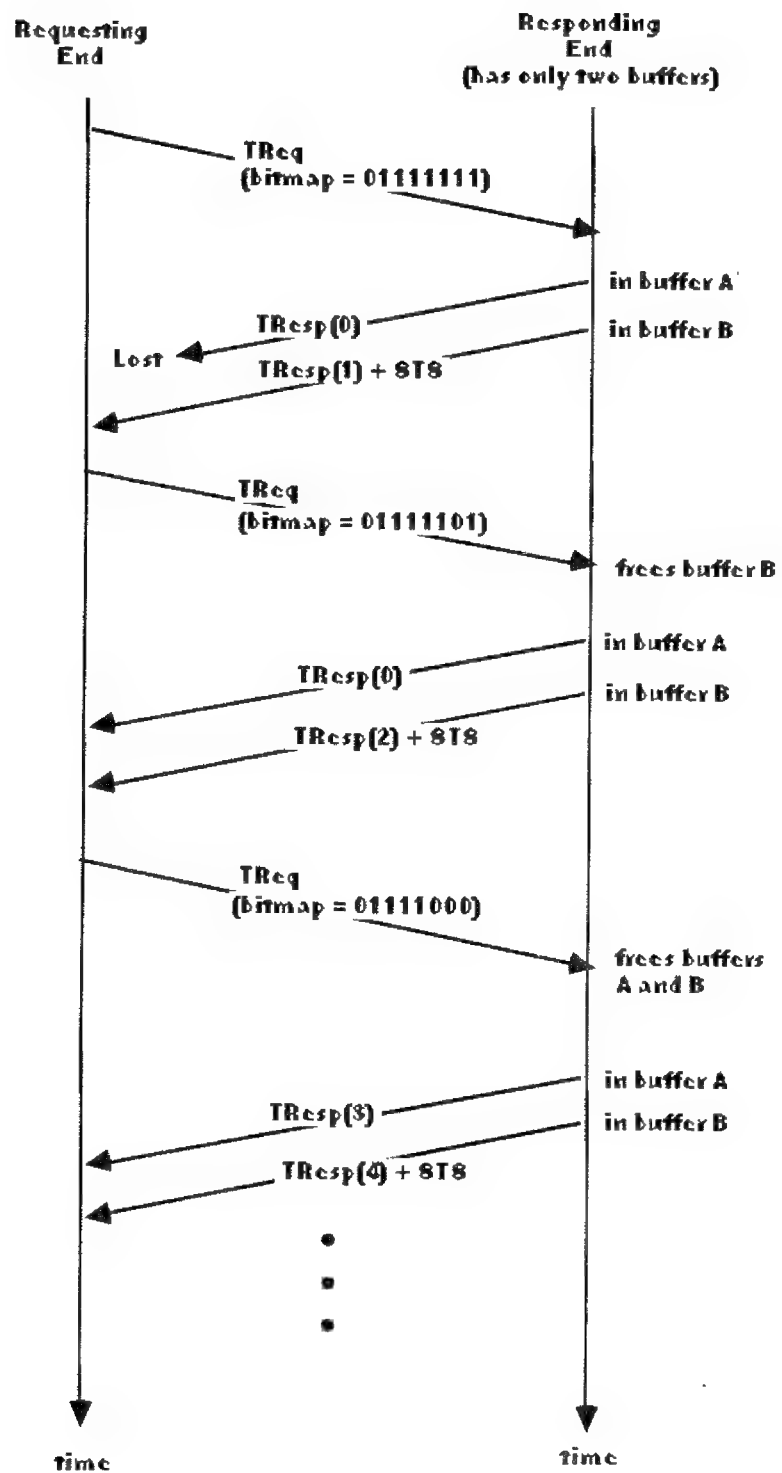


Figure VII-6: Use of STS




Figure VII-7: ATP Packet Format

VIII. Zone Information Protocol (ZIP)

About Zone Information Protocol

The Name Binding Protocol introduced the concept of a zone, defining a zone as an arbitrary subset of networks in the internet in which name lookups were performed. A given network was said to be in one and only one zone. It was specified in that chapter that the bridges were responsible for maintaining the mapping between networks and zone names, and for converting nodes' *broadcast requests* (BrRq's) into zone-wide lookups. This network-to-zone-name mapping is maintained by the Zone Information Protocol (ZIP). This chapter describes that protocol and its uses.

ZIP Services

An important feature of the Zone Information Protocol is that most of its services are transparent to the normal (non-bridge) node, as they are embedded in the use of the zone field of NBP. Unless a normal node wishes to be aware of the details of the zone structure of the internet, it does not have to implement a ZIP process in any form. Those normal nodes desiring this information may implement a small subset of the ZIP in their machines. The majority of ZIP, however, is implemented in the bridges. ZIP provides two major services: (1) the maintenance of the network-number-to-zone-name mapping of the internet, and (2) support for the various housekeeping commands that normal nodes may wish to implement for obtaining and possibly changing this mapping. The first portion of ZIP pertains only to bridges; the second concerns both bridges and normal nodes.

The Network-Number to Zone-Name Mapping

The Zone Information Table

Under stable conditions, each bridge maintains a complete network *number to zone* name mapping of the internet, known as the *Zone Information Table*, or ZIT. This table consists of one entry for each network in the internet. A ZIT entry is a tuple of the form <network number, zone name>. The zone name field can be NIL, indicating that the zone name of that network is currently unknown (this is a temporary condition); otherwise it is a (case insensitive) string specifying the zone name of that network.

The ZIP process in a bridge learns about new networks on the internet by monitoring RTMP's routing table. When it discovers an entry in the routing table that is not in the ZIT, it creates a new ZIT entry for that network with a zone name of NIL and initiates an attempt to discover its zone name. Likewise if the ZIP process discovers that the ZIT contains an entry whose network number is not in the routing table, it concludes that that network is no longer on the internet and removes its entry from the ZIT.

The Zone Information Socket: ZIP Queries and Replies

Associated with ZIP implementations in bridges is a statically assigned socket known as the *Zones Information Socket* or ZIS. This socket, socket 6, is opened by ZIP during its initialization. It is the socket to which other bridges address requests for zone information and through which a given bridge responds to those requests. These requests, known as *ZIP queries*, contain a list of network numbers that the requesting node is attempting to

determine the zone names of. A bridge receiving such a request responds with a *ZIP reply* listing the requested zone names of which it is aware (it does not respond if it is not aware of any of the requested zone names).

ZIT Maintenance

ZIP requires an additional field in the port descriptor of each bridge port: this is the zone name of that port's network. This field is only necessary for ports connected to AppleTalk networks, and furthermore only for those ports for which the bridge contains seed information. It is valid for this field to be NIL, indicating *unknown*, as long as the restrictions specified under "Zone Name Assignment" are adhered to.

Upon being turned on, each bridge's ZIT consists of one entry for each of its directly-connected AppleTalks for which it is a seed bridge. Both the network number and zone name fields of these entries are taken from the port descriptor for those networks. Additionally, ZIP monitors the routing table for the addition of any networks that are not in the ZIT (that is, networks of which the bridge has just become aware). When it discovers such a network, it creates a new ZIT entry, with a zone name field of NIL.

Whenever a ZIT entry is created with a zone name of NIL (either as just described, or at initialization time through being taken from a port descriptor), ZIP sends out a zone query in an attempt to discover the zone name of that network. This query is sent to the ZIS in the node specified as follows: the routing table entry for the given network is examined. If the next bridge field is zero (indicating the network is a directly-connected one), the query is broadcast; otherwise it is sent to the indicated bridge. The next bridge field specifies, in some sense, the start of a route to the desired network. The zone information for a given network can be viewed as propagating outward from that network. In contrast to RTMP however, the information is proactively requested as opposed to reactively received.

Bridges receiving a ZIP query asking for a zone name they are aware of respond (to the requesting socket) with a ZIP reply indicating that zone name. Bridges unaware of this information do not respond at all. As responses are received, the requesting bridge enters the specified zone names into the appropriate ZIT entries. For some queries, no reply may be received (the query may have been lost or the queried bridge may not yet have the information). ZIP does nothing in this case. ZIP does, however, maintain a background timer for the purpose of retransmitting these requests. Whenever this timer expires, ZIP retransmits one query for each entry in its ZIT whose zone name is still NIL. This query is transmitted in exactly the same manner as it was sent originally. In this way, zone names will propagate dynamically outward from the named network itself -- those bridges one hop away will receive the information on their first query, those two hops away may not receive it until their second, etc. Eventually on a stable internet, the ZIT in every bridge will contain the entire *network to zone* name mapping, and there will be no further ZIP activity.

ZIP also monitors the routing table to determine if a network goes down (that is, a network listed in the ZIT is not in the routing table). In this case ZIP deletes the corresponding ZIT entry (the network may come back up later with a different zone name, which must be re-discovered).

Changing Zone Names

Under stable conditions, each network's zone name is replicated in the ZIT of each bridge on the internet. Changing a network's zone name requires making sure that that piece of

information is changed in every bridge. Somehow each bridge must be notified of the change, because on a stable internet bridges are no longer sending out zone queries. One possible notification method would be an internet-wide broadcast of the change request. Internet-wide broadcasting, however, is not supported by DDP, and is both complicated and expensive (in terms of network traffic) to implement. Thus ZIP uses an alternate method.

The method used by ZIP for changing the zone name of a given network is conceptually a simple one. ZIP combines with RTMP to logically *bring down* the network. Once the network is down, its entry disappears from all the routing tables in all the bridges on the internet. As the network disappears from all the routing tables, ZIP also removes it from all the ZITs. Once this process is complete, ZIP combines with RTMP to bring the network back up with the new zone name. As a new network coming up, its zone name then propagates out into the internet through the normal zone query process. The network becomes known by the new zone name.

An important aspect of the zone name changing processes is that the network is only logically "down." It does not have to be brought down physically (although it can be if desired). More importantly, ZIP and RTMP combine in such a way as to enable internet traffic to still be routed *through* the "down" network. Although no traffic can be routed *to* the network, all other internet traffic flows normally through it. Thus even a network through which critical internet traffic flows can have its zone name changed with no effect on that traffic.

As alluded to in the RTMP chapter, ZIP requires an additional service of RTMP to enable it to change the zone name of a network: ZIP must be able to cause RTMP to remove an entry for one of its directly connected networks from its routing table, and to re-insert that entry at a later point. Given this service, ZIP can perform the name change. Details follow.

The request to change a network's zone name must originate from a node on that network. The node broadcasts a special *ZIP takedown* request to the ZIS. As a broadcast, it should be sent several times. All bridges on the network receive this request through their ZIS. Upon receiving the request, the ZIP process in each of these bridges instructs its associated RTMP process to delete the routing table entry associated with the network through which the takedown request is received. Since the routing information for a given network originates in and is "refreshed" by the bridges connected to that network, removing this information from all those bridges effectively brings the network down. Through the operation of RTMP, the entry for that network will disappear from the routing table of every bridge on the internet. The network is now "down." Bridges connected to that net can still forward packets between themselves, however, and thus packets can be forwarded through the net.

Once the network is completely down, it can be brought back up with a new zone name. This is done by broadcasting a *ZIP bringup* request to the ZIS on the net which is to be brought back up. This request specifies the new zone name the network is to be brought up with. The ZIP process in all bridges on the network receives the request and instructs its associated RTMP process to recreate the routing table entry for that network. The ZIP process also creates a ZIT entry for that network with the specified zone name. The network is now back up, with the new zone name. The fact that it is up propagates out through the internet via RTMP, and the new zone name propagates outward through the ZIP query process.

A network's zone name can of course also be changed by physically isolating the network from the internet, re-initializing all the bridges with the new zone name, and returning the

network to the internet. However, just as in the takedown/bringup case, the network can not be returned to the internet (brought back up) until information about it has disappeared from all the ZITs in the internet. This is an important point: in either case, a network can not be brought back up with a new zone name for a given amount of time after it is brought down. Although this parameter is somewhat a function of internet topology, ZIP defines it as a constant known as the *ZIP bringback time*. The exact value of the ZIP bringback time is defined in the section on parameters.

In general, non-bridge nodes need not be concerned with implementing these commands. Although it will be a non-bridge node which issues the ZIP takedown and ZIP bringup commands, it is envisioned that an application program will be provided which performs the following functions: issues the ZIP takedown command, waits for the ZIP bringback time, and issues a ZIP bringup command with the new zone name. Thus only the application program will have to implement these ZIP commands.

Note that ZIP provides no means of renaming an entire zone -- each network on the zone must be renamed individually.

Listing Zone Names

ZIP provides ways for normal nodes to obtain information about the zone structure of the internet. For one, normal nodes are free to issue ZIP queries to bridges. Through ZIP queries, normal nodes can obtain the zone name associated with any network on the internet, including their own. Since ZIP queries and responses are delivered only on a "best effort" basis, the normal node will also have to implement a timeout-and-retry mechanism to guarantee a response. In addition, the ZIP query mechanism provides no simple way of obtaining a complete list of all the zone names on the internet. For these reasons ZIP provides additional commands through which normal nodes can obtain zone information.

Obtaining zone information is a typical request/response type transaction. The normal node requests some information from a bridge, and the bridge responds with that information. For this reason, ATP is used for these commands. However, to prevent it from being necessary to implement a full ATP responder in all bridges, the following restrictions apply to these requests: (1) they are of the "atleast once" type, (2) they are short enough to fit entirely in the ATP user bytes, and (3) they ask for only a single response packet.

Two ATP requests are provided by ZIP. They are both sent to the ZIS in any bridge on the local network. The ATP retry count and interval are left up to the particular implementation.

The *GetZoneList* command is used to obtain a list of all the zones on the internet. Since all the zone names may not fit in one ATP response packet, the request contains an index at which to start including names in the response (zone names in the bridge are assumed to be indexed from one on up). To obtain the complete zone list, a non-bridge node sends a series of ATP requests. The first of these requests specifies an index of one. The response user bytes contain the number of zone names in that response packet, and whether there are any more zones whose names didn't fit in the response. If there are, the node should send out another request, with the index equal to the index sent in the last request plus the number of names in the last response. In this way a node can obtain the complete zone list.

It is an important feature of this command that zone names do not cross response boundaries. It is also important that, if more than one request is necessary in order to obtain the complete zone list, then all requests must be sent to the same bridge. This is because different bridges may have their zone lists stored in different orders. Note that a zero byte response will be returned by a bridge if the index specified in the request is greater than the index of the last zone in the list (and the user bytes will indicate "no more zones"). Although during periods when zone names are being changed, a station may receive a particular zone's name more than once, a bridge should make every attempt to send a particular zone's name only once in response to a *GetZoneList* command.

The *GetMyZone* command is used to obtain the name of the zone in which the requesting node resides (i.e. what '*' is equivalent to). This command is essentially a simplified *GetZoneList* request, where only one zone name is returned. A zero byte response will be returned if the bridge does not currently know the zone name of the zone (a temporary condition).

A bridge which is in the process of changing the zone name corresponding to a network to which it is directly connected, is not expected to respond correctly to a *GetMyZone* command from a node on that network. This restriction applies only to the period during which the zone name is being changed.

Packet Formats

Figure VIII-1 summarizes the ZIP packet formats. A ZIP packet is always sent with a short DDP header (i.e. on the local network). Queries are always sent to the ZIS in a bridge (or broadcast); replies are always sent from the bridge to the socket the query came from. Takedowns and Bringups are broadcast to the ZIS. The DDP type field in these four packets is set to 6 to indicate Zone Information Protocol. All four contain a header made up of a command byte (1=query, 2=reply, 3=takedown, 4=bringup) and a network number count (n), which should be equal to one in takedowns and bringups. Queries then contain n network numbers for which zone names are desired. Replies contain n network-number/zone-name pairs, with zone names being preceded by a length byte. Takedowns contain no data part; Bringups contain an unused network number field, and then the new zone name string.

The *GetZoneList* request contains a command code of 8 indicating "GetZoneList" and the desired starting index, both in the ATP user bytes. The response contains, again in the user bytes, a flag which is non-zero if this response contains the last zones in the zone list, and the number of zones contained in the data part. The *GetMyZone* request contains a command code of 7 in the user bytes (all other user bytes should be zero). The response will generally indicate one zone name in the user bytes, and contain that zone name.

NBP Routing in Bridges

As indicated in the NBP chapter, bridges contain an NBP process which is responsible for the conversion of BrRq requests to a zone-wide broadcast or LkUp requests. The process is a straightforward one. The bridge obtains, from its ZIT, a list of all the networks in the desired zone. It then uses DDP to send a directed-broadcast LkUp request to each of those networks. This directed broadcast is a normal DDP packet (generally with a long header), with a DDP destination node of \$FF. If the bridge is directly connected to the desired network, the packet should be broadcast on that net (LAP destination also \$FF, and a short

DDP header can be used). If not, the packet is sent to the "next bridge" field in the routing table entry for that network, which will forward the packet appropriately.

One DDP directed broadcast is sent for each of the networks in the zone. The DDP data part is exactly the same as that from the BrRq, except the command field is changed to LkUp. Note that NBP is defined such that the bridge's NBP process does not participate in any way in the responding process (the response is sent directly to the requestor through DDP). Also note that in addition to BrRq requests, bridges will also receive normal LkUp's on the NIS. These should not be forwarded.

Zone Name Assignment

RTMP defined the structure of a bridge's *port descriptor*. ZIP adds a zone name field to the port descriptor definition. Zone names are set into the port descriptors of bridge ports, and are then dynamically propagated out through ZIP to all bridges on the internet.

Only certain bridge ports need have zone names associated with them. These are the ports connected to AppleTalk networks. For each AppleTalk network, at least one bridge on that network must be configured with the network's zone name; all other bridges could have a zone name of NIL in their port descriptor associated with that net. If more than one bridge is configured with a non-NIL zone name for a given network, these names must be the same. Furthermore, only bridges which are seed bridges for purposes of specifying the network number should contain non-NIL zone names (i.e. if a bridge is not a seed bridge for routing information it should not be a seed bridge for zone information either).

Note that when the zone name of a network changes through a ZIP bringup command, all bridges on that net which have a non-NIL name in that net's port descriptor should be sure to change that name to the new one.

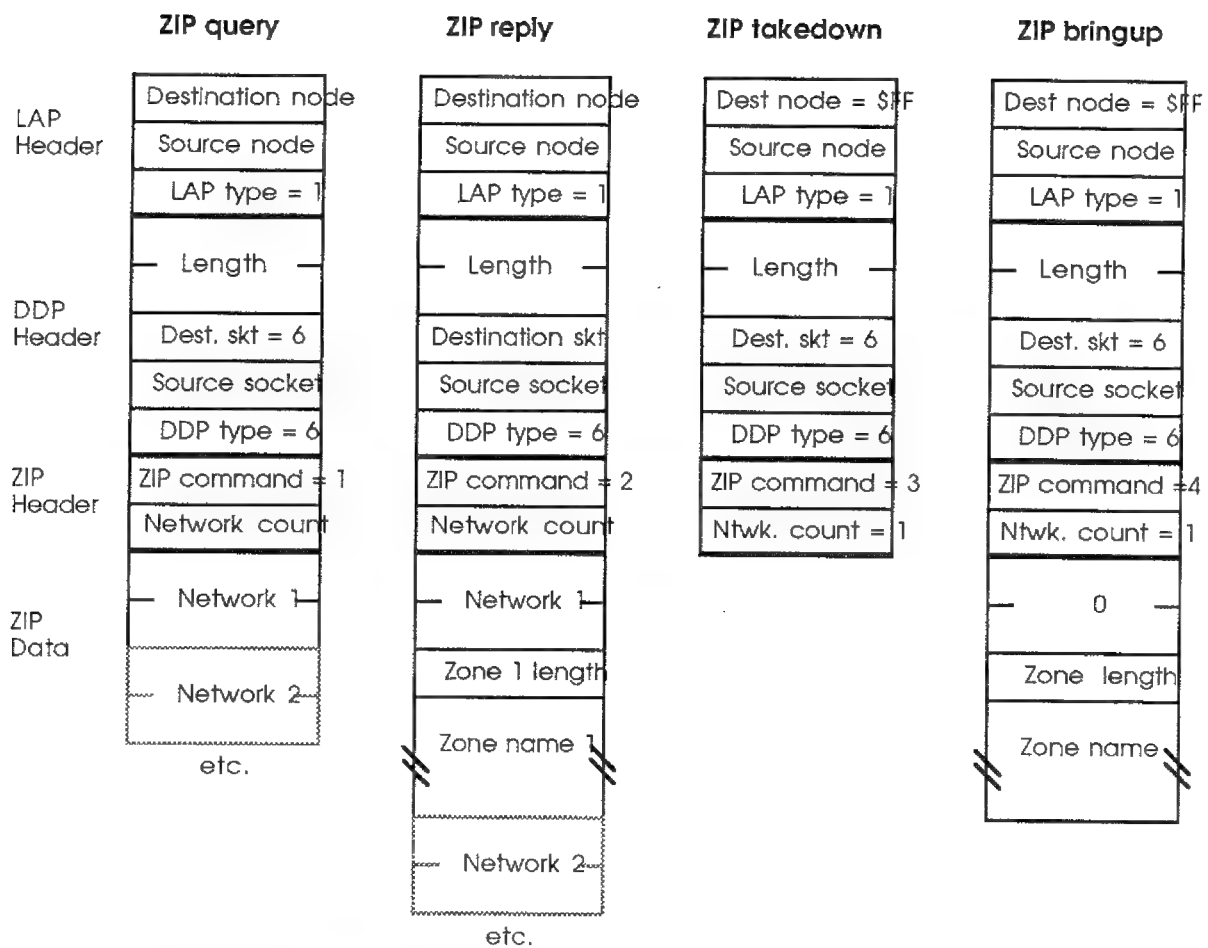
Timer Values

There are two parameter values associated with ZIP which must be specified. The first of these is the value for the query retransmission timer. This is defined as equal to the Send-RTMP timer, or ten seconds.

The second parameter is the ZIP bringback time -- the minimum time required between bringing a network down and bringing it back up with a new zone name. Since this is a constant regardless of internet topology, the worst-case internet must be used in determining it. This value has been somewhat conservatively defined since changing a network's zone name will be a very rare event. The value for the ZIP bringback time is ten minutes.

Implementation Notes

Although ZIP and RTMP are two separate and distinct protocols, changes in the ZIT are directly related to changes in RTMP's routing table. For this reason, bridge implementations may wish to combine the ZIT and routing table into one joint data structure, and the ZIP and RTMP processes into one joint module. This is perfectly acceptable.



ZIP GetZoneList or GetMyZone

to ZIS (6) in ABridge,
DDP type = "ATP" (3)

GetZoneList or GetMyZone Reply

to requestor,
DDP type = "ATP" (3)




Figure VIII-1. ZIP packet formats

IX. Printer Access Protocol (PAP)

About Printer Access Protocol

The AppleTalk Printer Access Protocol (PAP) is a session level protocol intended for use between workstations and print servers. PAP is a connection-oriented protocol which handles connection setup, maintenance and teardown in addition to data transfer. Multiple connections are allowed at both the workstation and server ends.

PAP envisions a print server node as containing one or more processes, referred to in this document as *servers*, which are accessed by workstations through PAP. Each of these server processes makes itself visible over the network by opening a service listener socket (SL socket) on which the corresponding server registers its name(s).

A PAP client in a workstation issues a PAPOpen call which initiates a connection-establishment dialogue with a server. The client specifies the server by its complete name; to process the PAPOpen call, PAP itself calls NBP to obtain the address of the server's listener socket from the server's name. PAP allows implementations in which the workstation's PAP client performs the NBP lookup directly and then makes the PAPOpen call with the address of the server's service listener socket.

Once a connection has been opened to the server, the PAP client at either end of the connection can receive data from the other end by issuing PAPRead calls, and can write data to the other end through PAPWrite calls. PAP uses ATP transactions (in exactly-once mode) to transfer the data.

When the data transfer has been completed a PAPClose call is issued by the PAP client on either end to close the connection.

At any time, the PAP client in the workstation can issue a PAPStatus call to determine the status of a server. PAP does not restrict the syntactic or semantic structure of this status information beyond specifying that, in the PAP packets, it is a string of at most 255 bytes preceded by a length byte. Note that the PAPStatus call can be issued even before a connection has been set up to the server.

PAP is not a symmetric protocol. There are several PAP calls for use only in server nodes. The first of these is the SLInit call. This is issued by each server (in a server node) when it is first started up, after it has completed its initialization. The SLInit call opens a service listener socket (an ATP responding socket) in the server node, and causes the server's name to be registered on that socket through NBP. Multiple SLInit calls may be issued to the PAP in the same server node -- each SLInit call opens a new service listener socket and registers the server name provided in the call on that socket.

A second call is used by the PAP client in the server to indicate to the node's PAP connection arbitration code that this client is ready to accept a new connection through a particular service listener socket opened via a prior SLInit call. This GetNextJob call primes the connection arbitration code to accept another connection establishment request from a workstation.

A SLClose call can be issued by a server client to close a given service listener socket and shut down any active connections for that client.

Two calls, PAPRegName and PAPRemName, can be used by a PAP client in the server node to register and remove (deregister) a server name for a specified service listener socket. This might be necessary for giving the server (on a particular service listener socket) more than one name, or to change the server name (for example, at server setup time) associated with a particular service listener socket.

One of the situations that PAP must deal with is the well-known case of half open connections. Such a connection is said to exist when one of the connection ends *dies* (or terminates the connection without informing the other end). Half open connections must be detected and torn down/closed. For this purpose, PAP maintains a connection timeout (at each end). Furthermore, each end of an open connection must send *tickling* packets to the other end on a periodic basis. The purpose of these packets is to inform the other end that the sender's end is open and *alive*. The receipt of any packet on a connection resets the connection timer at the receiving end. If the connection timer expires (that is, no packets have been received since the timer was last reset) then the decision is made that the other end is dead and the connection end is torn down.

A detailed discussion of the Printer Access Protocol (PAP) its interaction with (that is, use of) NBP and ATP, and the PAP client interface follows.

The Protocol

The basic model of a PAP-based server is that it processes jobs from workstations, where at any given time it can be processing a specific maximum number of jobs (this number is server implementation dependent and is not a PAP parameter). While the server is processing this maximum number of jobs, requests for initiating further jobs are not accepted (the requesting workstation is informed that the server is busy). While a server is processing a job, a connection is said to be *open* between the workstation being served and the server. There is a one-to-one correspondence between open connections and jobs being processed by the server. When the server is done with a particular job, the corresponding connection is closed and the server may, at its discretion, notify its PAP that it is able to accept another request for service from any workstation.

When a server process is first started, it goes through its internal initialization and then issues a SLInit call to its PAP code. This causes PAP to call ATP to open an ATP responding socket. This is the service listener (SL) socket for that server. Then, PAP calls NBP to register the server's name(s) and bind them to the SL socket. An ATP Receive-a-Request call is then issued by PAP on this socket (to allow the server can respond to PAPOpen or PAPStatus request packets). But the server itself (a PAP client in the server node) may still not be ready to accept a job/connection. PAP will still not accept connection opening requests through this SL socket. The server is said to be in a BLOCKED state. (see Figure IX-1).

After the SLInit call completes, the server process issues a series of GetNextJob calls to indicate that the server is ready to accept jobs. One such call is issued for each job it can accept at the time. The server is now in the WAITING state and is ready to accept jobs/connections.

PAP can support multiple servers within one node. Each of these servers is made available on a unique SL socket set up through a corresponding SLInit call.

PAP uses NBP to name and find a server's SL socket. Apart from these operations, all packets exchanged by PAP are sent through ATP, that is, they are ATP request or response

packets. Each such PAP packet contains in the ATP user bytes a one-byte quantity that indicates its PAP-type. Thus, for instance reference will be made to an ATP transaction of PAP-type OpenConn, etc.

Connection Establishment (Opening) Phase

A connection is a logical relationship between two PAP code entities (one in the workstation and the other in the server node). Data can be exchanged by two PAP clients only after a connection has been established/opened. Since PAP uses ATP to transfer data, the two communicating PAPs must, in the connection establishment phase, discover the address of the ATP responding socket of the other connection end. Also, the amount of data that can be transferred in an ATP transaction is of a maximum size equal to the available receive buffers at the end issuing the read requests. This maximum size (called the *flow quantum*) is sent by each end to the other in the connection establishment phase.

Connection establishment is initiated by a PAP client in a workstation by issuing a PAPOpen call. Such a client provides as a call parameter the complete name of the server. The PAP code obtains the complete internet address of the server's SL socket by issuing an NBP Lookup call. It then opens an ATP responding socket R_w , generates an 8-bit connection identifier ConnID and then sends a transaction request (TReq), with PAP-type OpenConn, to the server's SL socket. This packet contains the ConnID, the address of socket R_w , the flow quantum for the workstation, and a wait time used by the server for arbitration (discussed later).

All packets related to this connection (sent by either end) must contain this connection identifier. Packets with different connection ID's received through the sockets associated with the connection should be ignored by PAP. The workstation should be sure to generate the connection ID's in such a way as to minimize the likelihood of any two connections opened by that workstation having the same ID (especially connections established at about the same time).

When an ATP TReq of PAP-type OpenConn is received at the server's SL socket, PAP executes a connection acceptance algorithm (see Figure IX-1). If the server is BLOCKED (that is, there are no outstanding GetNextJob calls), then its PAP responds to the *OpenConn* with an ATP response of PAP-type OpenConnReply indicating *server busy*. Included in the OpenConnReply is a status string which is passed to the client and can be used for further detail on the busy state.

If however, the server is in the WAITING state (that is, there are one or more pending GetNextJob calls), then upon receiving an OpenConn (the first one since the server went into the WAITING state), its PAP goes into an arbitration (ARB) state for a fixed amount of time (two seconds). In the ARB state, the PAP receives all incoming "OpenConns" and tries to find the ones corresponding to the workstations that have been waiting for a connection for the longest amount of time. The idea is to implement a fairness scheme that accepts the requests generated by these stations over those from more recent entrants to the contest.

The time, in seconds, for which a station has been waiting (called the *WaitTime*) is sent with the OpenConn. When the first OpenConn is received since the server went to the WAITING state, the WaitTime value from that request is loaded into a variable associated with one of the pending GetNextJob calls. This GetNextJob call is marked as being tagged. During the ARB interval, whenever a new OpenConn request is received, the server examines all the pending GetNextJob calls to see if any one of them is not tagged. If

such a free pending GetNextJob call is found, then the WaitTime of the just received OpenConn is saved with that GetNextJob which is then marked as tagged.

If no free GetNextJob call is found among the pending calls, then PAP compares the just received OpenConn's WaitTime with the values saved in all the tagged pending GetNextJobs. If it is less than all of them, PAP responds to the just-received request with an OpenConnReply indicating *server busy*. If it is greater than one, PAP associates the new request with the GetNextJob with the smallest saved WaitTime, replacing that WaitTime with the one from the new request.

At the end of the ARB interval, PAP opens ATP responding sockets R_s for each connection request still associated with a GetNextJob and sends ATP responses of PAP-type OpenConnReply indicating *connection accepted* to the selected (but still pending) requests. These carry the ConnID received in the *OpenConn*, the address of socket R_s and the flow quantum of the server end (the flow quantum value for the server end is set by the SLInit call issued when the server is initialized). The corresponding connections are now open, and the jobs from the corresponding workstations can now be processed.

At the end of the ARB state, if there are no pending GetNextJobs then the server goes to the BLOCKED state, else it goes to the UNBLOCKED state. In the BLOCKED state there are no pending GetNextJob calls and the server cannot accept incoming OpenConn requests. However, if it is in the UNBLOCKED state then there are pending GetNextJob calls and the server can still accept additional connections/jobs.

Note that if the server is in the UNBLOCKED state, it has just been through the ARB state and has already opened connections to all workstations that have been waiting for a connection. Thus in the UNBLOCKED state, upon receiving the next OpenConn request it is not necessary to go into the ARB state -- as long as it has pending GetNextJob calls, the server accepts all incoming OpenConn requests and sets up connections immediately. As soon as it runs out of pending GetNextJob calls, it goes to the BLOCKED state until a GetNextJob is issued -- at which time it must go into the WAITING state.

At the workstation end, if a response of PAP-type OpenConnReply is received indicating that the server is busy (that is, in the BLOCKED state), then PAP waits some time (around 2 seconds) and issues another connection opening transaction. Each time it repeats this process it updates a "wait time" -- the time in seconds that it has been trying to open the connection. The current value of this WaitTime is sent with each OpenConn. Each of these OpenConn ATP transaction requests is issued with a retry count of 5 and retry interval of 2 seconds. The workstation's PAP should provide some way for its client to abort a PAPOpen request, but should otherwise keep trying until the connection is opened.

Data Transfer Phase

Once a connection has been opened PAP's data transfer phase is initiated. In this phase, PAP has two functions: to actually transfer data over the connection, and to detect and tear down half-open connections.

The detection of half-open connections is done by maintaining a connection timer (of duration 2 minutes) at each end of a connection. This timer is started as soon as the connection is opened. Whenever a packet of any sort is received from the other end of the connection, the timer is reset. If the timer expires (clearly, without receiving a packet from the other end) the connection is torn down. The presumption is made that the other end has

"died", has closed its connection or has become otherwise unreachable (for instance if an internet has become partitioned).

For this to work properly, it is important that even though no data is being exchanged on the connection, PAP exchanges control packets to signal that the connection ends are alive. This process is referred to as *tickling* and the control packets are called *tickling packets*. For this purpose, as soon as a connection is established, each end starts an ATP transaction with PAP type Tickle. This transaction, known as the *Tickle* transaction, has a retry count of infinity and a retry time interval of half the connection timeout. The *Tickles* must be ATP transactions of at-least-once type. Tickle packets are sent to the other end's ATP responding socket (that is, R_s or R_w). The receiver of such a packet must reset its connection timer but must not send a transaction response. These *Tickle* transactions are cancelled by each end when the connection is closed.

The basic data transfer model used by PAP is *read-driven*. When the PAP client at either end of the connection wishes to read data from the other end, it issues a PAPRead call. This call provides PAP with a read buffer (of size equal to this end's flow quantum) into which the data is to be read. As a consequence of this call, PAP calls ATP to send an ATP transaction request with PAP-type SendData, and the ATP bitmap reflecting the size of the call's read buffer. This transaction is issued with a retry count of *infinite* and a retry time interval of 15 seconds. It is sent to the other end's ATP responding socket. To prevent duplicate delivery of data to PAP's clients, all ATP transactions for the transfer of data use ATP's exactly-once mode and a sequence number. This is described in detail in the section "Duplicate Filtration."

The receipt of an ATP TReq packet with PAP-type SendData implies that there is a pending PAPRead at the other end. This *send credit* can be remembered by the PAP code, and used to service any pending or future PAPWrite calls issued by its client.

When a PAP client (at either end) issues a PAPWrite call, PAP examines its internal data structures to see if it has received a *send credit*. If it has, then it takes the data from the PAPWrite and sends it in ATP response packets with at most 512 bytes of ATP data each and of PAP-type Data (the EOM-bit is set in the last of these ATP response packets). If no send credit has been received, then PAP queues the PAPWrite call and awaits a *send credit* from the other end (that is, the receipt of an ATP request of PAP-type SendData from the other end). The amount of data to be sent in a PAPWrite call cannot exceed the flow quantum of the other end (PAPWrite calls that violate this restriction return immediately with an error message).

When a PAP client issues the last PAPWrite call for a particular job, it should ask PAP to send an End-of-File (EOF) indication with that call's data. The EOF indication is delivered to the PAP client at the other end (as part of the received information for a PAPRead call); this indication notifies the client that the other end is through sending data on this connection. Note that for this purpose the client may issue a PAPWrite call with no data to be sent; in this case, just an EOF indication is conveyed to the client at the other end.

Duplicate Filtration

As described in the ATP chapter, in the case of internets ATP exactly-once mode does not guarantee exactly-once delivery of requests -- it only guarantees that if a duplicate request is delivered to an ATP client, it can be ignored because all responses to it have been successfully received by the other side. PAP uses a sequence number in read (SendData)

requests to enable it to detect these duplicates and ignore them. Because PAP maintains only one outstanding read request at a time, this can be done in quite a simple way.

All read requests contain a sequence number in the last two ATP user bytes. The sequence number starts at 1 with the first read and takes on successive values up to 65535 before wrapping back around to 1 again. The value zero is reserved to mean *unsequenced*. Any SendData request received with a sequence number of zero should be accepted by PAP without duplicate checking. This is for compatibility with previous versions of PAP. If the sequence number is non-zero, PAP should verify that the sequence number is equal to one more than the sequence number of the last SendData request received. If this is not the case, the packet should be ignored as a duplicate of a previous request. Each side of the PAP connection must independently maintain both a sequence number for its SendData requests and a sequence number for the last SendData request received from the other side.

Connection Termination (Closing) Phase

When the PAP client at either end issues a PAPClose call, PAP closes the connection. Typically, after the workstation's PAP client has completed sending all data to the server and received an EOF in return, it will issue the PAPClose call. An ATP transaction request is sent to the other end with PAP-type *CloseConn*. An end receiving a *CloseConn* should immediately send back, as a courtesy, an ATP transaction response of PAP-type *CloseConnReply*. To close a connection's end, it is important to cancel any pending ATP transactions issued by that end, including the *tickle* transaction. An end receiving a *CloseConn* must cancel its pending ATP transactions for that connection as soon as it is able to do so.

At the server end, the receipt of the *CloseConn* will cause the connection to be torn down, but the server may continue to process the data pertaining to the job. When this is completed, the PAP client in the server may if it wishes to accept another job issue a *GetNextJob* call. Note that in fact the server may issue a *GetNextJob* call at any time in order to signal to its PAP code its willingness to accept another job. These *GetNextJob* calls are queued up by PAP and used to accept incoming *OpenConn* requests as discussed above in the Connection Establishment Phase.

A server may also close all open connections by issuing an *SLClose* call, which deregisters all names and closes its SL socket.

Status Gathering

Additionally, PAP supports status querying of the server through the *PAPStatus* call. It is not necessary to have a connection opened to the server to issue this call -- a workstation client can issue this call at any time. The call results in a *SendStatus* request packet being sent to the server specified in the call (the server can be specified by name, in which case PAP will call NBP to determine the server's address). The request is sent to the server's SL socket. The server's PAP responds with a *StatusReply* packet which contains a Pascal-format string (length byte first) specifying the server's status. This is done without passing the request to the PAP-client in the server. The PAP client in the server must have previously provided the status string to PAP through a *SLInit* or *HeresStatus* call. The *HeresStatus* call, which is implementation dependent, should be made by the PAP server client whenever the printer's status changes. Note that this status string is also returned by PAP in *OpenConnReply* packets.

PAP Packet Formats

PAP uses both NBP and ATP. The use of NBP is strictly for the purpose of registering a name on the server's SL socket and, given a server's name, for determining the address of its SL socket.

Packets sent by ATP in response to PAP calls include a PAP header. This is built using the user bytes of the ATP header, and in some cases by sending four or more bytes of PAP header in the data part of the ATP packet.

In all cases, the first of the ATP user bytes is the ConnID (except for SendStatus requests and Status replies for which this byte must be equal to zero). The second ATP user byte is the PAP-type of the packet. A list of PAP-type values follows in the next section. For packets of PAP-type equal to Data, the third ATP user byte is the EOF indication (non-zero indicates end-of-file). For packets of type SendData, the third and fourth bytes are the sequence number (high byte first). OpenConn and OpenConnReply packets contain, as part of the ATP data, the ATP responding socket numbers and the flow quantum to be used for the connection. The OpenConn request additionally contains the WaitTime; the OpenConn reply contains the open result and the status string. StatusReplies contain just the status string.

Figure IX-2 illustrates the PAP header for the different types of PAP packets.

PAP Type Values

The permissible PAP Type field values are:

```
OpenConn = 1
OpenConnReply = 2
SendData = 3
Data = 4
Tickle = 5
CloseConn = 6
CloseConnReply = 7
SendStatus = 8
StatusReply = 9
```

Values returned in the error code field of a OpenConnReply are:

```
NoError = 0          { No error - connection opened }
PrinterBusy = $FFFF  { Printer busy }
```

The Client-PAP Interface

In this section we take each PAP call and list the parameters the client must pass and the significant interface-level aspects of each call.

PAPOpen

This call is issued by a PAP client in a workstation to start/open a connection to a specified server.

- Call parameters:

- an indication as to the printer to which a connection should be opened (entity name or printer's SL socket address)
 - flow quantum (number of 512 byte buffers for reads)
 - a buffer in which the open status string is to be returned

- Returned parameters:

- result code
 - connection refnum

The open status string should be returned by PAP at any time it is received in a OpenConnReply, not just upon call completion. The client must use the connection refnum in order to refer to this connection in subsequent calls.

PAPClose

This call is issued by the PAP client in a workstation to close the connection specified by its connection refnum.

- Call parameters:

- connection refnum

- Returned parameters:

- result code

PAPRead

This call is issued by the PAP client at either end to read data from the other end over the connection specified by the connection refnum.

- Call parameters:

- connection refnum
 - buffer into which to read the reply

- Returned parameters:

- result code
 - size of the data read
 - the end-of-file indication

It is important to note that PAP assumes that the buffer into which the reply is to be read is equal to (no smaller than) the flow quantum specified in the PAPOpen call.

PAPWrite

This call is issued by the PAP client at either end to write data to the other end over the connection specified by the connection refnum.

- Call parameters:

- connection refnum
 - buffer with the data to be written
 - size of the data to be written
 - end-of-file indication

- Returned parameters:

- result code

If the data size is bigger than the flow quantum of the other end (value received during the connection establishment phase) the call will return with an error.

PAPStatus

This call is used by a PAP client in the workstation to determine the current status of the print server. It can be used at any time regardless of whether a connection has been opened by the PAP client to the server or not. Upon completion this call returns a string with the status message sent by the print server.

- Call parameters:

- an indication as to the printer from which status is being requested (entity name or printer's SL socket address)
 - a buffer in which the status string is to be returned

- Returned parameters:

- result code

In addition to these calls, the server uses the following four calls:

SLInit

This call is issued by the PAP client in the server to open a service listening socket and to register the server's name on this service listening socket. The client can also provide an initial status string.

- Call parameters:

- the entity name of the printer
 - the flow quantum for all connections to the printer (number of 512 byte buffers)
 - a status string

- Returned parameters:

- result code
 - the server refnum for the given printer

This server refnum must be used by the server when issuing subsequent GetNextJob calls in order to identify the listener socket for which the GetNextJob call is being made. Thus the PAP code in the server node must return a unique server refnum for each SInit call.

GetNextJob

This call is issued by the PAP client in the server whenever it is ready to accept a new job through the listener socket specified by the server refnum.

- Call parameters:

server refnum

- Returned parameters:

result code

a refnum to refer to this connection by in subsequent calls (read/write/close)

SLClose

This call is issued by the PAP client in the server whenever it wants to close down a given server process.

- Call parameters:

server refnum

- Returned parameters:

result code

PAPRegName

This call is used in server nodes only. It registers a name (as an entity name for the print server) on the server's listening socket corresponding to the specified server refnum.

- Call parameters:

server refnum
server name to register

- Returned parameters:

result code

PAPRemName

This call is used in server nodes only. It deregisters a name from the server's listening socket corresponding to the specified server refnum.

- Call parameters:

- server refnum
 - server name to deregister

- Returned parameters:

- result code

HeresStatus

This call is issued by the PAP client in the server to provide PAP with a new status string. This call should be issued any time the status string has changed.

- Call parameters:

- server refnum
 - status string

- Returned parameters:

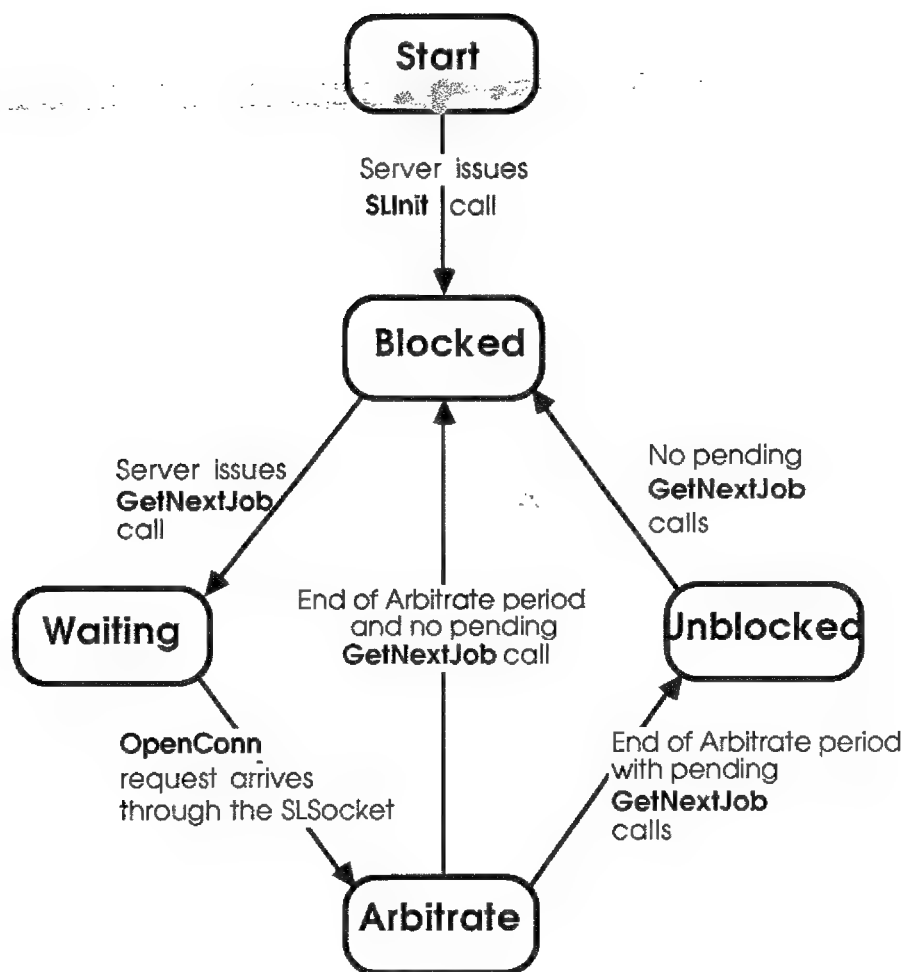
- result code

The Apple LaserWriter™

This section lists some of the specifics of the PAP-client implementation on the Apple LaserWriter™ printer.

The flow quantum used by the LaserWriter is 8.

The LaserWriter can handle only one job at a time, so it never has more than one GetNextJob outstanding. Essentially, the UNBLOCKED state does not exist on the LaserWriter; it is either WAITING, arbitrating (ARB), or BLOCKED (busy).



OpenConn requests received through the SL Socket are ignored if the state is BLOCKED.

OpenConn requests received through the SL Socket are serviced immediately when in the UNBLOCKED state, until no more GetNextJob calls are pending (then go to BLOCKED state).

Figure IX-1. Server State Diagram

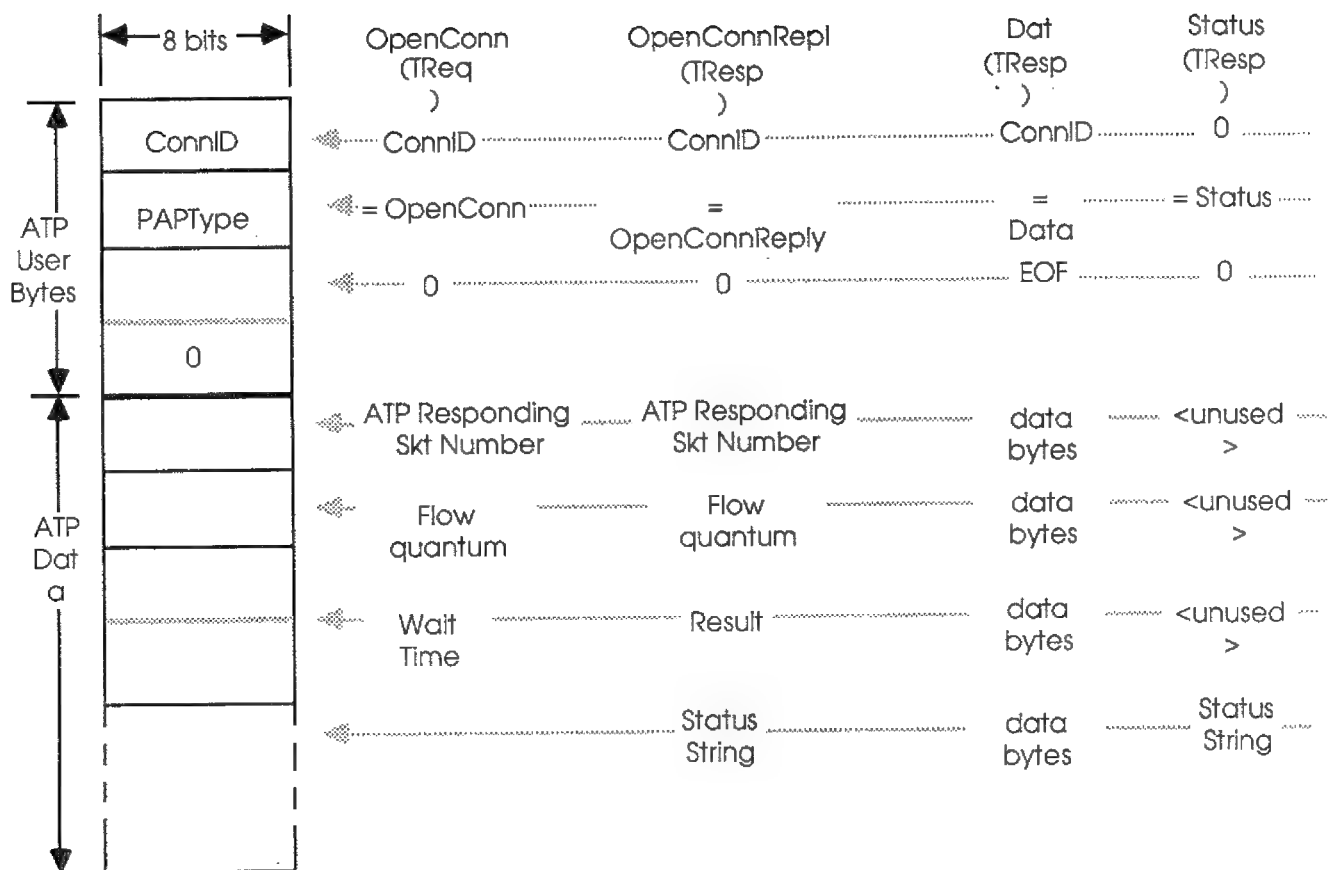


Figure IX-2. PAP packet format

X. Echo Protocol (EP)

About Echo Protocol

The Echo protocol is implemented on a particular node as a statically-assigned socket (socket number 4 known as the *Echoer* socket) and an Echoer process on this socket. The Echoer listens for packets received through this socket. Whenever a packet is received (see Figure X-1 for the packet format) the echoer examines its DDP protocol type and the size of the DDP data in this packet. If this protocol type field is not equal to 4 (the DDP protocol type for the Echo protocol) or the DDP data length is zero, then the Echoer discards the packet and ignores it. However, if the just received packet has a DDP protocol type equal to 4 and its DDP data length is one or more, then the Echoer examines the first byte of the DDP data (this is the Echo Protocol header). If this is equal to 1 (Echoer Function = Echo request), then the Echoer changes this field to 2 (Echoer Function = Echo reply) and calls DDP to send the packet back to the original sender.

A client process wishing to use the Echo protocol's service must first determine the internet address of the node it wishes to obtain the echo from (it will probably use NBP for this purpose). Then it calls DDP to send an Echo request packet to the Echoer socket in that node. The client can send the datagram through any socket it has opened (the Echo response will come back through this socket). The client then waits for the receipt of the Echo reply packet. Note that the client can set the Echo data part of the request packet to any pattern it desires and then examine the data in the reply (which will be the data sent in the request). This may be used by the client, for instance, to distinguish among the replies to various Echo requests it may have sent.

There are several situations in which an Echo reply is never received by the client. The Echo protocol packets could get lost in the network system. The target node may not have an *Echoer*. The target node may be unreachable or down. The client must use prudence in determining how long to wait for the echo reply before concluding that one of these situations exists. It might choose to try several times before concluding that the remote node will not respond at all.

This simple protocol can be used for two important purposes. In the first place, it can be used by any DDP client to determine if a particular node, known to have an Echoer, is accessible over an internet. A more significant use is to obtain an estimate of the round trip time for a typical packet to a remote node, usually a server. This is very useful in developing client-dependent heuristics for estimating the timeouts to be specified by clients of ATP, ASP and other higher level protocols. This is, in general, a difficult problem in the case of internets.

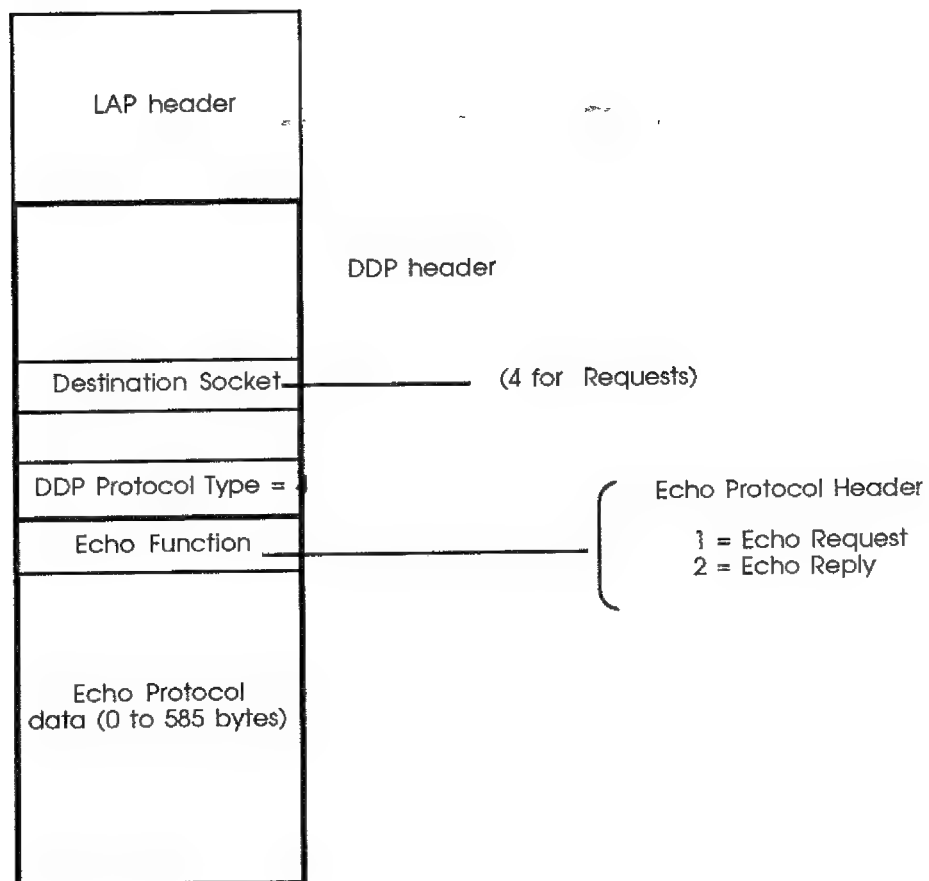


Figure X-1. Echo Protocol packet

XI. AppleTalk Session Protocol (ASP)

This chapter specifies version 1.0 of the AppleTalk Session Protocol. It supersedes the preliminary draft document version 1.0/3. The protocol as specified herein is the final version 1.0 of this protocol. This protocol was developed in joint work by Apple Computer and Centram Systems West.

About AppleTalk Session Protocol

A general conceptual model encompassing a wide variety of higher-level network services consists of a workstation issuing a sequence of commands to a server that executes these commands and returns the corresponding results to the workstation. An important example of such a service is a filing service through which file system commands can be conveyed to a file server for their execution.

At the transport level, the AppleTalk architecture provides a reliable transaction service via the AppleTalk Transaction Protocol (ATP) that can be used by such higher-level services for the conveyance of commands to servers. Nevertheless, ATP does not provide the full range of transport services needed by such higher-level protocols. In this document we describe the *AppleTalk Session Protocol* (ASP) designed specifically for the use of these higher-level protocols.

ASP is a client of ATP; it adds value to ATP to provide the level of transport service needed for higher-level workstation-to-server interaction.

Sessions, Commands, and Command Replies

Central to ASP is the concept of a *session*. Two network entities, one in a workstation and the other in a server can set up an ASP session between themselves. This is a logical relationship identified by a unique *session identifier*. For the duration of the session, the workstation entity can, through ASP, send a sequence of *commands* on the session to the server entity. ASP ensures that the commands are delivered in the same order as they were sent and conveys the results of these commands (known as a *command reply* or simply a *reply*) back to the workstation entity.

It is important to note that ASP sessions are inherently asymmetrical. The process of setting up a session is always initiated by the workstation entity (when it wishes to use the server entity's advertised service). Once the session is set up, the workstation client of ASP sends commands on the session and the server client of ASP replies to the commands. ASP does not allow the server client of ASP to send commands to the workstation client. It does, however, provide an *attention* mechanism by which the server can inform the workstation that it is in need of attention.

Although ASP guarantees that commands issued by the workstation end of a session are delivered to its server end in the same order as they were issued, it is beyond the scope of ASP to ensure that the commands are in fact executed and completed in that order by the server end. This is clearly the responsibility of the ASP client at the server end.

Many workstation entities can have sessions to the same server entity at the same time. ASP uses the session identifier to distinguish between commands received on these various

sessions. The session identifier of a session is unique among all the sessions to the same server.

What ASP Does Not Do

ASP does not in any way understand the syntax or the semantics of the commands sent by its workstation clients on a session. Commands are conveyed basically as blocks of bytes whose interpretation is left to the server end/client of the session. Likewise, the command replies are sent back over the session to the workstation client without any syntactic/semantic understanding/interpretation by ASP.

It has been an important goal in the design of ASP to make its client interface independent of the lower-level transport protocols. This makes the higher-level clients of ASP transportable from one network to another with a minimum of modification.

As an important consequence it is necessary to separate out from ASP both the mechanism by which a server entity advertises its service and the manner in which workstation entities look for this advertised service. How this is realized depends intimately on the transport and naming mechanisms of a particular network; ASP assumes that these functions are taken care of by the ASP clients themselves and are outside its scope.

For instance, let us discuss the mechanism used on AppleTalk. A server entity wishing to advertise its service on the AppleTalk network calls the AppleTalk Transaction Protocol (ATP) to open an ATP responding socket known as the *Session Listening Socket* (SLS). The server entity then calls the Name Binding Protocol (NBP) to register a name on this socket. [At this point, the server entity calls ASP, via the *SPInit* call discussed later, to pass it the address of the SLS. Then, ASP starts listening on the SLS for session opening requests coming in over AppleTalk.] Now, a workstation entity wishing to access this advertised service uses NBP to discover the address of the SLS (this is known as the *Entity Identifier*). Then, it calls the workstation ASP to open a session by sending a session opening packet to the SLS. [Note that the operations of setting up an SLS and of looking for the SLS via NBP are done outside the scope of ASP. ASP's participation starts with the process of setting up a session.]

ASP does not provide a user authentication mechanism. This is expected to be the responsibility of the higher-level protocol being used by the clients of ASP.

ASP does not provide a mechanism to allow the use of a particular session by more than one server entity. Such multiplexing of a session can be done by the ASP clients if the command identifiers are broken into ranges by the higher-level protocols and managed completely outside the scope of ASP. We recommend the use of separate sessions to access different services on the same server.

ASP Services and Features

ASP provides the following services to its clients:

- setting up (*opening*) and tearing down (*closing*) sessions
- sending *commands* (on an open session) to the server and returning its *command replies* (which might include a block of data)
- *writing* blocks of data from the workstation to the server end of the session
- sending *attentions* from the server to the workstation

In addition, the workstation client can ask ASP to obtain *service status information* from the SLS without opening a session. The status information (provided by the server entity) is not semantically intelligible to ASP.

Opening and Closing Sessions

Before any sessions are opened, both workstation and server ASP clients should interrogate ASP to discover the maximum sizes of command and replies allowed by the underlying transport mechanism. This information may be used by both ends to determine whether or not the underlying transport services are adequate to their needs, or just to optimize their commands and replies.

After discovering the network address of the desired service's SLS (the entity identifier), the workstation client calls ASP to open a session to this service (the client passes the entity identifier in this call). As a result of this call, ASP sends a special *OpenSession* packet (an ATP request) to the SLS; this packet carries the address of a workstation socket to which session maintenance packets (discussed later) are to be sent. This socket will be referred to as the *workstation session socket* (WSS). If the server entity is unable to set up a session, it returns an error (in the ATP response packet known as an *OpenSessReply*), otherwise it returns a session acceptance indication together with a *session identifier*, and the number of the server-end socket of the session (this socket will be referred to as the *Server Session Socket* or SSS). In all further communication over this session all packets must carry this *SessionID* and must be sent to the SSS.

ASP allows protocol version number verification in this session opening dialogue. The ASP in the workstation sends an *ASP protocol version number* in the *OpenSession* packet. This identifies the version of the ASP protocol that the workstation is using. If the server's ASP is unable to handle this version it will return an error in the *OpenSessReply* and the session will not be opened.

A session can be closed by the ASP client at either end. This is done by issuing the appropriate command to that end's ASP which duly informs the other end and then immediately tears down the session. If the session termination was initiated by the workstation client, then a session termination notification is sent to the SSS (this is an ATP request which carries the *SessionID*). If it was initiated by the server client, then the notification is sent to the WSS, again as an ATP request carrying the *SessionID*.

Whenever a session is terminated, the ASP clients at both ends must be notified of this occurrence so appropriate higher-level actions may be taken. This is easily done at the server end because it is generally listening for incoming commands on the session. But at the workstation end (if the server end closed the session) this is done the next time that the ASP client tries to issue a command on that session. The actions taken by an ASP client, upon being informed of the closing of a session, could be of various sorts depending on the higher-level function. For instance, the server end might choose to free up resources allocated for that session. If the higher-level service is a filing service, it might decide to flush all files, etc., opened on that session.

Session Maintenance

A session will remain open until it is explicitly terminated by the ASP client at either end or until one of the session's ends *dies* or becomes unreachable. ASP includes a mechanism

known as *session tickling* which is initiated as soon as a session is opened. The idea is that each end of the session periodically sends a packet (known as a *tickle* packet -- an ATP request) to the WSS or the SLS to inform the other end that it is alive and well. If either end fails to receive *any* packets on a session for a certain predefined *session maintenance time-out* it assumes that the other end has *died* or has otherwise become unreachable. At that time the session is closed.

A simple way for an ASP session end to generate tickle packets is to issue an ATP request with its retry count set to infinite. This ATP request must be cancelled when the session is closed.

Session Requests on an Open Session

Once a session has been opened, the workstation client of ASP can send a sequence of commands over the session to the server end. These commands are delivered in the same order as they were issued at the workstation end, and replies to the commands are returned to the workstation end by ASP. These commands can be classified into two types that differ in the direction of the primary flow of data:

1) ASP Commands: These are very similar in nature to ATP requests. The ASP workstation client sends a command (encoded in a variable size command block) to the server-end client requesting it to perform a particular function and send back a variable size command reply. Examples of such commands could vary from requests to open a particular file on a file server, to reading a certain range of bytes from an already opened file. In the first case a small amount of reply data is returned, in the second case a multi-packet reply might be generated. Each ASP Command translates into an ATP request sent to the SSS and the command reply is received as a possibly multi-packet ATP response (see Figure XI-3). In any case ASP does not interpret the command block or in any way participate in the command's function. It simply conveys the command block, encoded in a higher-level format unknown to ASP, to the server end of the session, and returns the command reply to the workstation-end client. The command reply consists of a four-byte *command result* and a variable size *command reply block*.

2) ASP Writes: In this case the ASP client in the workstation wishes to convey a variable size block of data (known as the *WriteData*) to the server end of a session and expects a reply. ASP uses ATP as its underlying transport protocol. Since ATP is a protocol in which a requesting end essentially *reads* a multi-packet block of data from the responding end, it is necessary to translate the ASP Write into possibly two transactions. First, ASP sends an ATP request to the SSS carrying the ASP Write's control information (this can be called the *Write command block*). (This is illustrated in Figure XI-5.) The server end examines this information and determines whether it wants to go ahead with *pulling* the data from the workstation end. If not, it returns an error in the ATP response packet (this error is conveyed to the workstation client as the four-byte *command result*). If it wants to pull the data, then the server end sends an ATP request to the WSS to *pull* the data from the workstation end. This latter request could generate a multi-packet ATP response carrying the write data to the server. The server end upon receiving the write data and performing the particular function requested in the ASP Write, then responds to the original ATP request with the appropriate error message (this error is conveyed to the workstation client as the four-byte *command result*).

Additionally, once a session has been opened, the server client can send an *attention request* to the workstation client. This is an ATP-ALO request sent to the WSS. The sole purpose of this request is to alert the workstation client as to the server's need for attention.

ASP delivers two bytes of *attention data* (from the request's ATP user bytes) to the workstation client and acknowledges the attention request (with an ATP response), but it is the workstation client's responsibility to act on the request. An example of the use of the attention mechanism might be if a server desires to notify a workstation as to a change in its status. Upon receiving the attention request, the workstation could then issue an ASP command to the server to find out the details of the status change.

Sequencing and Duplicate Filtration on Sessions

ASP ensures the delivery on a session of commands and writes to the server's end in the same order in which they were issued at the workstation end. This is done by including a *sequence number* in the appropriate packets exchanged by ASP.

The use of sequence numbers also allows ASP to add robustness to the ATP-XO service. Although ATP-XO will guarantee that a request is delivered to the ATP client exactly once if the source and destination nodes are on the same AppleTalk, this might not be the case over an AppleTalk internet. In the latter case it is possible for a copy of the ATP Request that was delayed in a bridge node to be delivered as a duplicate after the original transaction has been completed. Thus a duplicate transaction will be delivered. This problem is inherent to transaction protocols and can be eliminated when the concept of a session, as with ASP, is introduced and the transactions belonging to a particular session carry a sequence number which can be used to filter out delayed duplicates.

Getting Service Status Information

ASP provides an out-of-band service to allow its workstation clients to obtain a block of service status information from the SLS without the need for opening a session. In the server the status block is provided to ASP by the corresponding ASP client and is returned in response to *ASPGetServerStatus* commands received at the SLS.

ASP Client Interface

Having provided the preceding overall description of ASP and its services, we can now proceed to a more detailed specification, first of the service interface offered by ASP to its clients at the workstation and server ends, and then of the internal packet formats and details of how ASP uses ATP.

ASP's service interface has been designed to make it as independent as possible of the underlying AppleTalk transport mechanisms. The primary motivation for this is to allow easy porting of the higher-level (ASP clients) protocols to networks other than AppleTalk and to simplify some of the problems in the design of internet gateways. Nevertheless, the internals of ASP are intimately related to the AppleTalk Transaction Protocol, and thus ASP itself is not directly portable to other networks.

Server Side

To start with, the server's ASP client should issue a *SPGetParms* call to retrieve the maximum values of command block size and quantum size. The *MaxCmdSize* is the maximum size command that can be sent to the server. The *QuantumSize* returned by this call is the maximum size reply that can be sent to a command or the maximum size of data

that can be transferred in a *Write* request. Since ASP is built on top of ATP the value of *MaxCmdSize* returned will be 578 bytes and *QuantumSize* will be 4624 bytes (8 ATP response packets with 578 data bytes each). For client-compatible session protocols implemented on other networks, these values will be different.

SPGetParms

Inputs:

none.

Outputs:

MaxCmdSize-- maximum size of a command block

QuantumSize -- maximum data size for a command reply or a write

Errors:

none

Before any workstation can open a session with a server, the server's SLS must be established. This must be done by the ASP client at the server end. This client must issue an *OPEN-ATP-Responding-Socket* call to create the SLS. Then it must call NBP to register the appropriate server entity name on this socket. At this point, the SLS is set up so that workstations can discover its network address through an NBP lookup. However, workstations still cannot open sessions with this service entity. For this purpose, the ASP client in the server must now issue an *SPInit* call to ASP (passing to ASP the network address, known as the *SLEntityIdentifier*), followed by one or more *SPGetSession* calls.

SPInit

Inputs:

SLEntityIdentifier -- SLS network identifier

ServiceStatusBlock -- block with status information

ServiceStatusBlockSize-- size of status information block

Outputs:

SPError -- error code returned by ASP

SLSRefNum -- reference number for the SLS

Errors:

TooManyClients -- ASP implementation cannot support another client

SizeErr -- *ServiceStatusBlockSize* is greater than *QuantumSize*

This call is issued by the ASP client after having opened and named the SLS. The call passes the (network dependent) *SLEntityIdentifier* to ASP as well as a *ServiceStatusBlock*. This block is used to hold the service status information to be returned in reply to *GetStatus* requests received at the SLS. The *SLEntityIdentifier* is the complete internet address of the SLS.

SPInit returns the *SLSRefNum* (this is unique among all SLS's on the same server node) that is used in the *SPGetSession* call to make reference to the SLS passed in the *SPInit* call.

SPGetSession

Inputs:

SLSRefNum -- reference number for the SLS

Outputs:

SPError -- error code returned by ASP

SessRefNum -- session reference number

Errors:

ParamErr -- unknown *SLSRefNum*

NoMoreSessions -- implementation cannot support another session

This call is issued by the ASP client to allow it to accept an *OpenSession* command received on the SLS identified by the *SLSRefNum*. Each *SPGetSession* request authorizes ASP to accept one more *OpenSession* request. The call completes when such a request is received on the SLS and a corresponding session has been opened. The *SessRefNum* is returned to the server ASP client and must be used in all further calls to ASP that refer to this session. Clearly the *SessRefNum* must be unique among all sessions open to the server.

SPCloseSession

Inputs:

SessRefNum -- session reference number

Outputs:

SPError -- error code returned by ASP

Errors:

ParamErr -- unknown *SessRefNum*.

This call is issued by the ASP client to close the session identified by *SessRefNum*. As a result of this call that value of the *SessRefNum* is invalidated and cannot be used in any further calls. Furthermore, all pending activity on the session is immediately canceled. ASP clients who want all pending activity to be completed before the session is closed, should make sure of this before making the *SPCloseSession* call.

SPGetRequest

Inputs:

SessRefNum -- session reference number
ReqBuff -- buffer for receiving the command block
ReqBuffSize -- buffer size

Outputs:

SPError -- error code returned by ASP
ReqRefNum -- request identifier
SPReqType -- SP level request type
ActRcvdReqLen -- actual size of the received request

Errors:

ParamErr -- unknown *SessRefNum*
BufTooSmall -- *ReqBuff* is too small to hold the entire command block
SessClosed -- session has been closed

After a session has been opened, the ASP client in the server must issue *SPGetRequest* calls to provide buffer space for the receipt of requests on that session. The size of the buffer for receiving the command block sent with the request depends on the higher-level protocol, but need be no greater than *QuantumSize*.

After a request has been received, this call completes and returns a unique request identifier *ReqRefNum* and a one-byte quantity that identifies the SP level type of request. The permissible values of *SPReqType* are *Command*, *Write*, and *CloseSession*. If the received command block does not fit in the *ReqBuff*, then ASP returns as much as will fit along with a *BufTooSmall* error.

When this call completes, the caller is given the size of the received command block in the parameter *ActRcvdReqLen*.

If the session times out and an *SPGetRequest* is pending, the call will complete with an *SPError* value of *SessClosed*. If no call is pending, the next *SPGetRequest* call issued on the session will complete immediately with an error.

SPCmdReply

Inputs:

SessRefNum -- session reference number
ReqRefNum -- request identifier
CmdResult -- four byte command result
CmdReplyData -- command reply data block
CmdReplyDataSize -- size of command reply data block

Outputs:

SPError -- error code returned by ASP

Errors:

ParamErr -- unknown *SessRefNum* or *ReqRefNum*;
bad *CmdReplyDataSize* (negative value)

SizeErr -- *CmdReplyDataSize* is greater than *QuantumSize*
SessClosed -- session has been closed

If the request returned by the *SPGetRequest* call has *SPReqType* = *Command* then the ASP client must respond to it with an *SPCmdReply* call to ASP. The value of *ReqRefNum* passed with this call must be exactly the same as that returned by the corresponding *SPGetRequest* call. Two items must be conveyed to the workstation end of the session: a four-byte *command result* and a variable size *command reply block*. The actual values, format and meaning of the *CmdResult* and of the *CmdReplyData* are transparent to ASP. Note that the *CmdReplyData* can not be of size greater than the *QuantumSize* returned by the *SPGetParms* call (that is, *CmdReplyDataSize* must be no greater than *QuantumSize*), or else a *SizeErr* will be returned and no *CmdReplyData* will be sent to the workstation.

SPWrtContinue

Inputs:

SessRefNum -- session reference number
ReqRefNum -- request identifier
Buffer -- buffer for receiving the data to be written
BufferSize -- size of the buffer

Outputs:

SPError -- error code returned by ASP
ActLenRcvd -- actual amount of data received into *Buffer*

Errors:

ParamErr -- unknown *SessRefNum* or *ReqRefNum*;
 bad *BufferSize* (negative value)
SessClosed -- session has been closed

If the request returned by the *SPGetRequest* call has *SPReqType* = *Write* then the ASP client must respond to it with either an *SPWrtContinue* or a *SPWrtReply* call to ASP. The value of *ReqRefNum* passed with these calls must be exactly the same as that returned by the corresponding *SPGetRequest* call.

Exactly how the ASP client decides which of these calls to make depends on the higher-level protocol, but here is a description of the general idea. Upon receiving a request of *SPReqType* = *Write*, the ASP client examines the command block received with the request. This should contain, in the format appropriate to the higher-level protocol, a description of the type and parameters of the higher-level write operation being requested. The ASP client should use this command block information to decide if it can successfully carry out the operation being requested. If the operation cannot be carried out it should issue the *SPWrtReply* call with the appropriate higher-level protocol value in *CmdResult* indicating its inability to execute the operation and the reason why. If however, the operation can be carried out, then the ASP client in the server should initiate the process of actually transferring the data to be written from the workstation end of the session. This is done by issuing the *SPWrtContinue* call. A *SPWrtReply* call should also be issued upon completion of the write.

For instance, the higher-level client could be a filing protocol requesting the ASP client in the server to write a certain number of bytes to a particular file. If no such file exists, the server end should send back a *no such file* indication by issuing a *SPWrtReply* call. Otherwise it issues an *SPWrtContinue* call with a buffer into which the write data can be brought from the workstation, followed by an *SPWrtReply* once it has finished the write request.

The maximum size of the write data that will be transferred is equal to the *QuantumSize*. Providing a larger buffer is of no use.

SPWrtReply

Inputs:

SessRefNum -- session reference number
ReqRefNum -- request identifier
CmdResult -- four byte command result
CmdReplyData -- command reply data block
CmdReplyDataSize -- size of command reply data block

Outputs:

SPError -- error code returned by ASP

Errors:

ParamErr -- unknown *SessRefNum* or *ReqRefNum*
 bad *CmdReplyDataSize* (negative value)
SizeErr -- *CmdReplyDataSize* is greater than *QuantumSize*
SessClosed -- session has been closed

This call is issued by the ASP client in the server in order to terminate either successfully or unsuccessfully a *Write* command received through *SPGetRequest*. With this call the ASP client provides ASP with the four-byte *CmdResult* and the variable size command reply data block *CmdReplyData* (at most *QuantumSize* bytes) to be conveyed to the workstation end client. If a *SizeErr* is returned, no *CmdReplyData* will be sent to the workstation.

SPNewStatus

Inputs:

SLSRefnum -- reference number for the SLS
ServiceStatusBlock -- block with status information
ServiceStatusBlockSize -- size of status information block

Outputs:

SPError -- error code returned by ASP

Errors:

ParamErr -- unknown *SLSRefnum*

SizeErr -- *ServiceStatusBlockSize* is greater than *QuantumSize*

This call is used by the ASP client to update the *ServiceStatusBlock* first supplied in the *SPInit* call. The previous status information is lost. All subsequent *SPGetStatus* calls issued by workstations will retrieve the new status block.

SPAttention

Inputs:

SessRefnum -- session reference number

AttentionCode -- two-byte attention code (must be non-zero)

Outputs:

SPError -- error code returned by ASP

Errors:

ParamErr -- unknown *SessRefnum* , zero *AttentionCode*

NoAck -- no acknowledgment from workstation side

This call sends the attention bytes to the workstation and waits for an acknowledgment.

Workstation Side

SPGetParms

Inputs:

none.

Outputs:

MaxCmdSize-- maximum size of a command block

QuantumSize -- maximum data size for a command reply or a write

Errors:

none.

This call is exactly the same as the *SPGetParms* call for the Server side

SPGetStatus

Inputs:

SLSEntityIdentifier -- SLS network identifier

StatusBuffer -- buffer for receiving the status information

StatusBufferSize -- size of this buffer

Outputs:

SPError -- error code returned by ASP
ActRcvdStatusLen -- size of status information received

Errors:

NoServer -- server not responding;
BufTooSmall -- *StatusBuffer* is not big enough to hold entire status

This call is used by a workstation ASP client to obtain status information corresponding to a particular SLS. If the status information received is too large to fit into the *StatusBuffer* provided with the call, then an appropriate *BufTooSmall* value of *SPError* is returned but as much of the status information as fits is put in the *StatusBuffer*.

SPOpenSession

Inputs:

SLSIdentifier -- SLS network identifier
AttnRoutine -- Attention routine indicator (implementation dependent)

Outputs:

SPError -- error code returned by ASP
SessRefNum -- session reference number

Errors:

NoServer -- server not responding
ServerBusy -- server cannot open another session
BadVersNum -- server cannot support the offered version number
NoMoreSessions -- no more sessions available at workstation

This call is issued by an ASP client after it has obtained the network address/identifier of the SLS via an NBP lookup. If a session is successfully opened, then a *SessRefNum* is returned to the caller and should be used on all subsequent calls referring to this session. If a session cannot be opened, for whatever reason, an appropriate *SPError* value is returned. *AttnRoutine* specifies, in an implementation-dependent manner, a routine to invoke upon receipt of an attention request from the server.

SPCloseSession

Inputs:

SessRefNum -- session reference number

Outputs:

SPError -- error code returned by ASP

Errors:

ParamErr -- unknown *SessRefNum*.

This call can be issued at any time by the ASP client to close a session previously opened via an *SPOpenSession* call. As a result of this call that value of the *SessRefNum* is invalidated and can not be used in any further calls. Furthermore, all pending activity on the session is immediately cancelled. ASP clients who want all pending activity to be completed before the session is closed, should make sure of this before making the *SPCloseSession* call.

SPCommand

Inputs:

SessRefNum -- session reference number
CmdBlock -- command block to be sent
CmdBlockSize -- size of command block
ReplyBuffer -- buffer for receiving the command reply data
ReplyBufferSize -- size of this buffer

Outputs:

SPError -- error code returned by ASP
CmdResult -- four byte command result
ActRcvdReplyLen -- actual length of command reply data received

Errors:

ParamErr -- unknown *SessRefNum*
SizeErr -- *CmdBlockSize* is larger than *MaxCmdSize*
SessClosed -- the session has been closed
BuffTooSmall -- the reply buffer cannot hold the whole reply

Once a session has been opened, the ASP client can send a command to the server end by issuing a *SPCommand* call to ASP. A command block of maximum size *MaxCmdSize* can be sent with the command. (This maximum command block size for ASP is 578 bytes). Thus, if *CmdBlockSize* is larger than this maximum allowable size, then the call completes with *SPError* equal to *SizeErr*; in this case, no effort is made to send anything out over the network to the server end.

In response to a command, the server end returns two quantities: a four-byte *CmdResult* and a variable length command reply which will be returned in the *ReplyBuffer*. The size of the command reply actually received is returned in *ActRcvdReplyLen*. Clearly, this can be no larger than *QuantumSize*, so it is possible that only part of the reply is returned in this call. It is the responsibility of the ASP workstation-end client to generate another command to retrieve the rest of the reply.

SPWrite

Inputs:

SessRefNum -- session reference number
CmdBlock -- command block to be sent

CmdBlockSize -- size of command block
WriteData -- data block to be written
WriteDataSize -- size of data block to be written
ReplyBuffer -- buffer for receiving the command reply data
ReplyBufferSize-- size of this buffer

Outputs:

SPError -- error code returned by ASP
CmdResult -- four byte command result
ActLenWritten -- actual number of bytes of data written
ActRcvdReplyLen -- actual length of command reply data received

Errors:

ParamErr -- unknown *SessRefNum*
SizeErr -- *CmdBlockSize* is larger than *MaxCmdSize*
SessClosed -- the session has been closed
BuffTooSmall-- the reply buffer cannot hold the whole reply

This call is made by the ASP client to write a block of data to the server end of the session. The call first delivers the *CmdBlock* (no larger than *MaxCmdSize*) to the server end client of ASP and as described above, then the server end can actually transfer the *WriteData* over or return an error (delivered in the *CmdResult*).

The actual amount of data sent over will be less than or equal to *WriteDataSize* and will in any case never be larger than *QuantumSize*. The amount of write data actually transferred is returned in *ActLenWritten*.

In response to a write, the server end returns two quantities: a four-byte *CmdResult* and a variable length command reply which will be returned in the *ReplyBuffer*. The size of the command reply actually received is returned in *ActRcvdReplyLen*. Note that this can be no larger than *QuantumSize*.

Packet Formats and ASP Internals

In this section we spell out the internal details of ASP including packet formats. This is done by discussing the major operations and for each case presenting the internal packet exchanges and their structure.

Opening a Session

When the workstation client issues a *SPOpenSession* call, ASP issues an ATP-XO transaction request addressed to the SLS (see Figure XI-1). This ATP transaction request packet is known as an ASP *OpenSess* packet. The server's ASP returns an ATP transaction response packet known as an *OpenSessReply* packet.

The *OpenSess* packet (see Figure XI-9) carries in its ASP header (contained entirely in the ATP UserBytes) a one-byte *SPCmdType* field = *OpenSess*, a one-byte field containing the WSS socket number, and a two-byte ASP version number field. For this version of the ASP protocol the version number field must be equal to hexadecimal \$0100.

Upon receiving an *OpenSess* packet (an ATP transaction request with the *SPCmdType* = *OpenSess*) on an SLS, the server's ASP checks to see if there is a *SPGetSession* pending on that SLS. If no such call is pending then it returns a *ServerBusy* error in the *OpenSessReply* packet and the session is not opened. If such a call is pending then ASP checks the ASP version number in the *OpenSess* packet. If unacceptable, a *BadVersNum* error is returned. If this is acceptable then ASP opens an *ATP-Responding-Socket* SSS and generates a unique (per SLS) one-byte session identifier *SessionID*. Then it creates its internal session management data structures in which the WSS socket number received in the *OpenSess* packet is saved together with the *SessionID*, the SLS socket number, etc. Then the *OpenSessReply* packet is sent back. This contains in its ASP header (contained entirely in the ASP UserBytes) a two-byte *ErrorCode* (returned to the client as *SPError*), the one-byte *SessionID*, and the socket number of the SSS. Now the server end of the session is active. At this time the tickling process is initiated at the server end (details discussed later).

The workstation's ASP upon receiving the *OpenSessReply* examines its *ErrorCode* field. If this indicates *noErr*, then the *SessionID* and the SSS are taken from the packet and together with other control information saved in a session management data structure. At this point, the workstation end of the session is active and the tickling process is initiated at the workstation end (details discussed later).

The session management data structure must contain the session's identifier, the socket number of the other end of the session (that is, the WSS or the SSS), and a two-byte *LastReqNum*. When the session is opened, the *LastReqNum* is initialized to zero.

Getting Service Status

Because a *SPGetStatus* call can be made and serviced without opening a session, the corresponding packets do not carry a *SessionID* nor do they have a sequence number field. The workstation's ASP issues an ATP-ALO transaction request addressed to the SLS. Thus an ATP request known as a *GetStat* packet (see Figure XI-2) is sent to the SLS. This has (see Figure XI-9) the *SPCmdType* = *GetStat* with the rest of the three ATP UserBytes being unused and thus set to zero.

The ASP at the server end upon receiving a *GetStat* packet (that is, *SPCmdType* = *GetStat*) returns as the possibly multi-packet ATP response up to eight *StatusReply* packets. Each of these has (see Figure XI-10) the four ATP UserBytes equal to zero. The status information block provided in the *SPInit* or *SPNewStatus* call is sent as the ATP data of the *StatusReply* packets. The status information is packed into the packets with as many bytes as will fit (that is, all but the last *StatusReply* packet will have 578 bytes of the status information in them).

ASP Commands

When the ASP client in the workstation makes an *SPCommand* call, then ASP sends an ATP-XO request to the SSS of the indicated session (see Figure XI-3). This *CmdReq* packet has (see Figure XI-10) the *SPCmdType* = *Cmd*, and the *SessID* of the session and a two-byte sequence number. The sequence number must be generated using the following algorithm:

```
If LastReqNum = 65536 then LastReqNum := 0 else
    LastReqNum := LastReqNum+1;
```

Sequence Number := LastReqNum;

In effect, the sequence number will be one greater than the sequence number of the last command sent on the session.

The *CmdBlock* provided in the *SPCommand* call is sent in the ATP data part of the *CmdReq* packet. Thus the *CmdBlock* cannot be larger than 578 bytes in size.

At the server end, the ASP upon receiving the *CmdReq* packet delivers it to the ASP client (if there was a pending *SPGetRequest*, otherwise the request is ignored). The ASP client in the server then makes a *SPCmdReply* call with which it passes to ASP a four-byte command result and a variable size command reply data block (of maximum size equal to the *QuantumSize*; for AppleTalk this is 4624 bytes). The ASP generates from one to eight ATP response packets that it sends back to the source of the *CmdReq* packet. These *CmdReply* packets have the four ATP UserBytes set to zero except for the first *CmdReply* which carries the command result in these UserBytes. The command reply data block is broken up into up to eight pieces which are sent in the ATP data part of these packets (see Figure XI-10) packing as many bytes as will fit in each packet (all but the last packet must contain 578 bytes of the command reply block).

ASP Writes

When the ASP client in the workstation makes a *SPWrite* call, ASP sends an ATP-XO request to the SSS of the indicated session (see Figure XI-5). This *WriteReq* packet has (see Figure XI-10) the *SPCmdType* = *Write*, and the *SessID* of the session and a two-byte sequence number. The sequence number must be generated using the algorithm discussed above.

The *CmdBlock* provided in the *SPWrite* call is sent in the ATP data part of the *WriteReq* packet. Thus the *CmdBlock* cannot be larger than 578 bytes in size.

At the server end, the ASP upon receiving the *WriteReq* packet delivers it to the ASP client (if there was a pending *SPGetRequest*, otherwise the request is ignored). The ASP client in the server determines if it can process the request, presumably by examining the contents of the command block.

If the ASP client in the server decides that it is unable to process the request, it encodes an appropriate higher level protocol error message in the four-byte command result and/or command reply data and makes a *SPWrtReply* call to ASP. Then as shown in Figure XI-5, an ATP response packet known as a *WriteReply* is sent back to the source of the *WriteReq* (see Figure XI-10 for its format).

If, however, the ASP client in the server decides that it can process the request, then it reserves a buffer for the data and makes an *SPWrtContinue* call to ASP. This causes the ASP in the server to send an ATP-XO transaction request to the WSS. This *WriteData* call (see Figure XI-11) carries the Session Identifier and the Sequence Number taken from the *Write* packet (used by ASP to match the *WriteData* with the corresponding *Write*). It has a two-byte ATP data field in which the size in bytes of the buffer reserved by the server client for the write data is inserted. The workstation then returns the data in the transaction response packets (*WriteDataReply* packets) to the *WriteData*. At that time the data is delivered to the server-end ASP client. The latter then issues an *SPWrtReply* call to ASP which causes it to send a *WriteReply* packet (this is the ATP response to the original *Write*; see Figure XI-10 for the format of the *WriteReply* packet).

Reply Size Error Checking

SPGetStatus, *SPCommand* and *SPWrite* all require a buffer into which to put the server's reply, and the length of that buffer. ASP must be able to verify that the entire reply fits into the specified buffer. In general, this information is available from ATP. The workstation-side ASP will always issue an ATP request specifying the number of responses expected as equal to the number of 578 byte blocks that will fit into the caller's buffer, plus one if there is any remainder (buffer size MOD 578 non-zero). There are two cases where size errors could occur that should be examined.

First, the server's response is of a size such that the number of response buffers the server-side ASP wishes to return is equal to the number requested, but the response still is too big to fit in the caller's buffer. In this case the workstation ASP can determine there was an error because ATP will indicate that the last response did not fit in the last response buffer.

Second, the server's response is of a size such that the number of response buffers the server-side ASP wishes to return is greater than the number requested by the workstation side. In this case, the server side should return as much of the response as it can (that is, completely fill each requested response, including the last one, with 578 bytes of data). As long as the caller's buffer size was not an exact multiple of 578 bytes, the workstation ASP can again determine a size error occurred because the last response again will not fit in the last response buffer.

However, if the caller's buffer size was an exact multiple of 578 bytes, ASP cannot determine, solely on the basis of information returned by ATP, if a size error occurred. If the last response buffer is completely filled with data, it could be because either the response was exactly the requested size, or the response was greater than the requested size but the server side had no way of passing this information back to the workstation. To be able to resolve this situation, *the workstation-side ASP must ask for one additional response in this case*. If the response size is less than or equal to the caller's buffer size, this response will not be returned (the EOM bit will be set). However, if the response size is greater than the caller's buffer size, this additional response will be returned, making the workstation-side ASP aware of the error. Note that if the caller's buffer is exactly *QuantumSize* big (which is a multiple of 578), this situation does not apply, because this is the most data that can possibly be sent by the server side.

Tickles and Session Maintenance

Tickle packets (these are ATP TReq packets with *SPCmdType* = *Tickle*) must be sent by each end while a session is open (see Figure XI-4). The tickle packets are sent by the workstation to the SLS and by the server to the WSS. Tickle packets contain in the ASP header (the ATP UserBytes) the one byte *SPCmdType* = *Tickle*, the *SessionID*, and two unused bytes (see Figure XI-9). Tickle packets are sent by starting an ATP transaction with retry count equal to infinite and time-out equal to 30 seconds.

The session maintenance at each end is done by starting a *session maintenance timeout* of 2 minutes. Whenever any packet (tickle or otherwise) is received on the session this timer is restarted. If the timer expires (that is, if no packet is received for 2 minutes) then it is concluded that the other end of the session has *died* or has become unreachable, and the session is closed.

Attention Requests

When the ASP client in the server makes a *SPAttention* call, ASP sends an ATP-ALO request to the WSS of the indicated session (see Figure XI-8). This *Attention* packet requests one response buffer and has (see Figure XI-9) the *SPCmdType = Attention*, the *SessID* of the session and one word (two bytes) of *attention data*. This attention data is passed by the server client to ASP to be delivered to the workstation client along with the attention request. It is uninterpreted by ASP except for the fact that ASP requires it to be non-zero.

The workstation-side ASP, upon receiving an *Attention* request, should immediately respond with an *AttentionReply*. This serves as an acknowledgment of the request, and completes the *SPAttention* call on the server side. The workstation ASP should then, in an implementation-dependent manner, alert its client as to the attention request and pass the client the attention data.

Closing a Session

When the ASP client at either end makes a *SPCloseSession* call, a session closing ATP-ALO transaction is initiated. (see Figures XI-6 and XI-7).

If the session closing was initiated by the workstation client, then the *CloseSess* packet (an ATP transaction request) is sent to the SSS in the server, where it is delivered to the ASP client as part of the *SPGetRequest* mechanism. The server's ASP generates the TResp -- a *CloseSessReply* packet -- without the ASP client's intervention. Immediately upon sending the *CloseSessReply* packet the server end of the session is closed, all pending activity including tickles is canceled at that end, and no further server end calls to this session are accepted. The workstation end, immediately upon receiving the *SPCloseSession* call, cancels all pending activity on the session including tickles, and does not accept any more calls from its ASP client referring to that session. The workstation end can choose (this is implementation dependent) to retransmit the *CloseSess* packet (remember, it's an ATP transaction request) several times; it will close its end of the session either as soon as the *CloseSessReply* is received or the retries expire.

It is possible that on the server side no *SPGetRequest* was pending when the close request was received. In this case, everything should proceed as above and the session should be marked as closed. The server client, however, cannot be informed until his next request for that session. At this time, ASP should return a *SessClosed* or other error.

If the session closing was initiated by the server end then the *CloseSess* is sent to the WSS. The workstation's ASP generates the *CloseSessReply* without client intervention. The session is closed immediately upon receipt of the *CloseSess* packet. The ASP client in the workstation is informed of the session being closed when it makes a call to ASP and refers to that session. The server end can choose (this is implementation dependent) to retransmit the *CloseSess* packet several times; it will close its end of the session either as soon as the *CloseSessReply* is received or the retries expire.

The formats of the *CloseSess* and *CloseSessReply* packets are given in Figure XI-9. Note that the *CloseSess* packet does not include a sequence number and hence must be accepted by the receiving end without sequence number verification. Also, its receipt should immediately lead to the cancellation of all pending activity on the session, including tickles.

CloseSess packets are sent simply as a courtesy to the other end, and to potentially speed up the process of closing a session at both ends. Note that it is possible that the *CloseSess* packet or the *CloseSessReply* sent by either end can get lost; then the end initiating the session closing activity (let's assume it has chosen to send the *CloseSess* packet only once) will not receive an acknowledgment to its request and will close the session and stop tickles when its retries expire. If the other end in fact did not receive the *CloseSess* packet at all, it will not know that the session has been closed. However, this half-open session will be detected when the still active end's session maintenance timer expires, at which time it will close its own end of the half-open session.

SPCmdType Values

The standard values of *SPCmdType* are as follows:

```

CloseSession = 1
Command = 2
GetStat = 3
OpenSess = 4
Tickle = 5
Write = 6
WriteData = 7
Attention = 8.

```

All other values are invalid in this field.

SPErrors Values

The standard values of *SPErrors* are as follows:

<u>Error</u>	<u>Hex Value</u>	<u>Decimal Value</u>	<u>Where Returned</u>
NoError	0	0	Both sides (*)
BadVersNum	\$FBD6	-1066	Workstation (*)
BufTooSmall	\$FBD5	-1067	Workstation
NoMoreSessions	\$FBD4	-1068	Both sides
NoServers	\$FBD3	-1069	Workstation
ParamErr	\$FBD2	-1070	Both sides
ServerBusy	\$FBD1	-1071	Workstation (*)
SessClosed	\$FBD0	-1072	Both sides
SizeErr	\$FBCF	-1073	Both sides
TooManyClients	\$FBCE	-1074	Server
NoAck	\$FBCE	-1075	Server

Values -1060 through -1065 are reserved for implementation-dependent errors. All other values are invalid in this field. Error codes marked with a (*) are the only ones actually transmitted through ATP (on the *OpenSession* call).

Timeouts and Retry Counts

ASP uses ATP transactions as its basic building blocks. For each of these transactions it is necessary to provide a retransmission timeout value and maximum retry count.

Most transactions used by ASP (except: opening a session, getting service status, requesting attention and closing a session) use a retry count of infinite. This involves no danger of leading to a deadlock, since half-open connections (other end *dead* or unreachable) are easily detected through the tickling mechanism.

The maximum number of retries used by the session opening, getting service information and attention transactions should be specifiable by the ASP user. How this is done is an implementation issue not discussed in this document.

The timeout value to be used in any of the transactions (with the exception of tickles) and how this is specified by the user or built into ASP is an implementation issue and is not specified in this document.




Figure XI-1. Session opening transaction




Figure XI-2. Get Status transaction




Figure XI-3. Command transaction




Figure XI-4. Tickling transactions

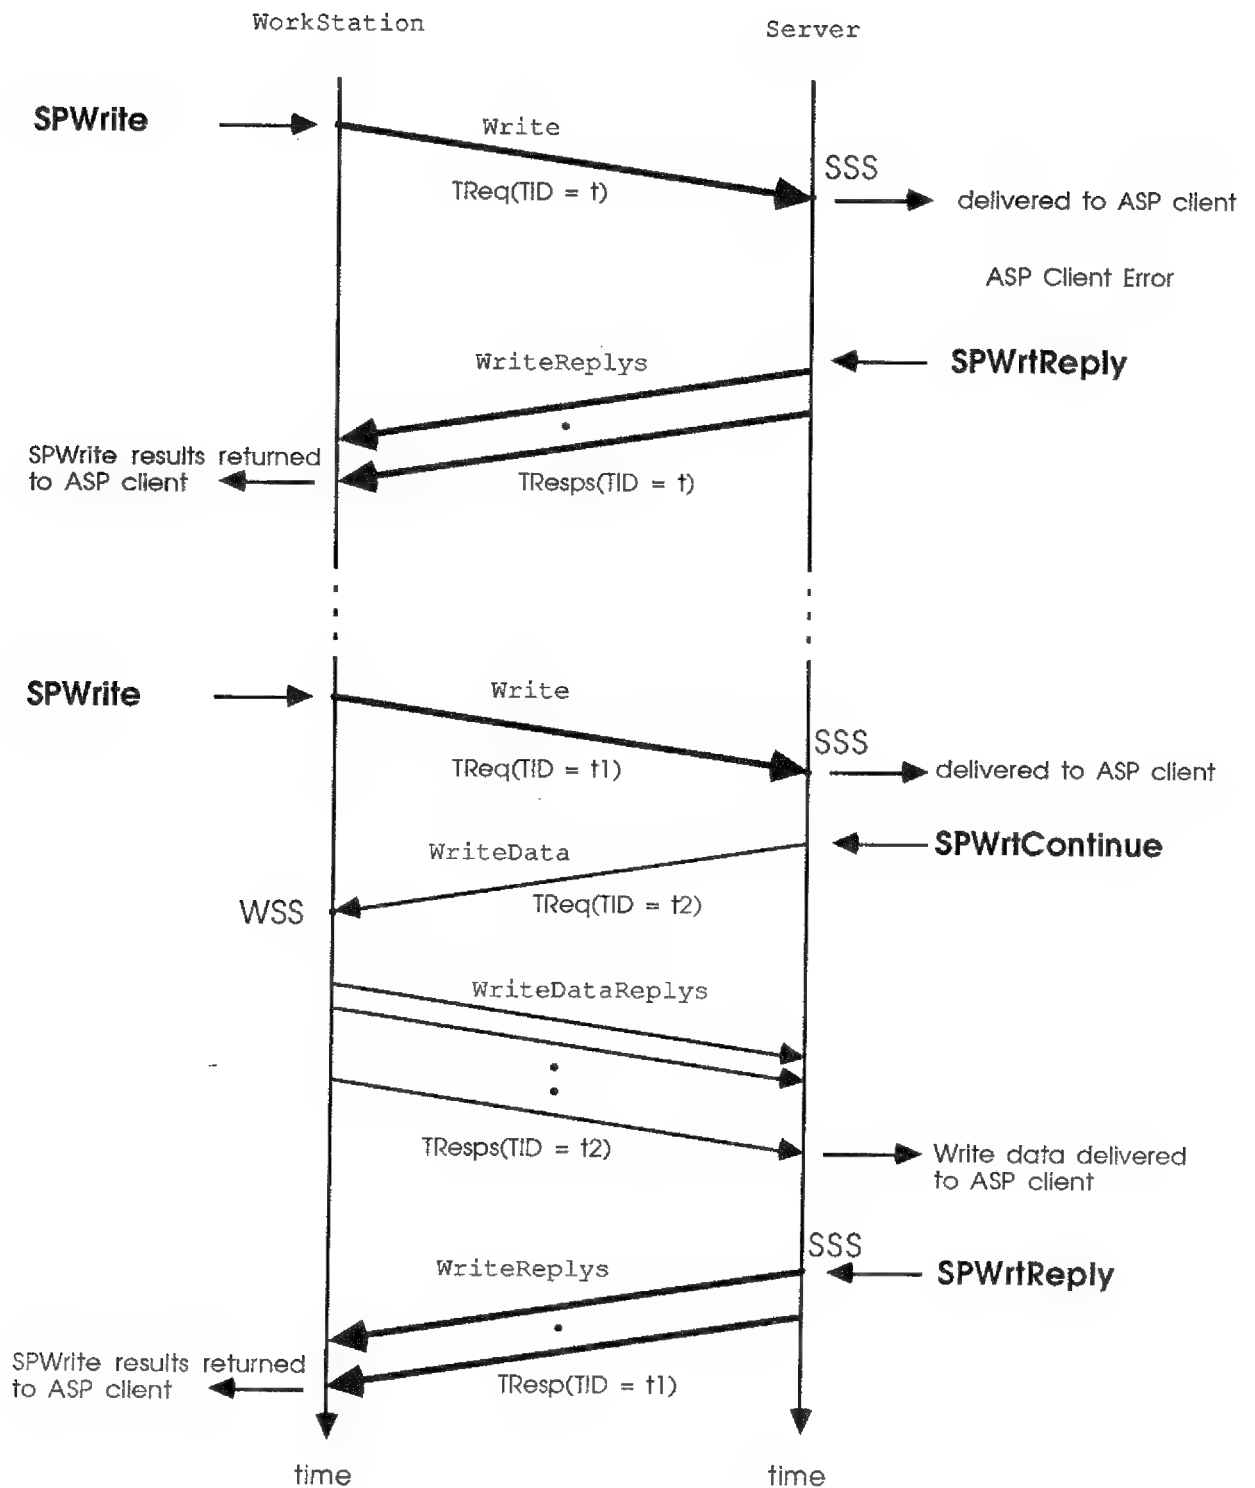


Figure XI-5. Write transaction

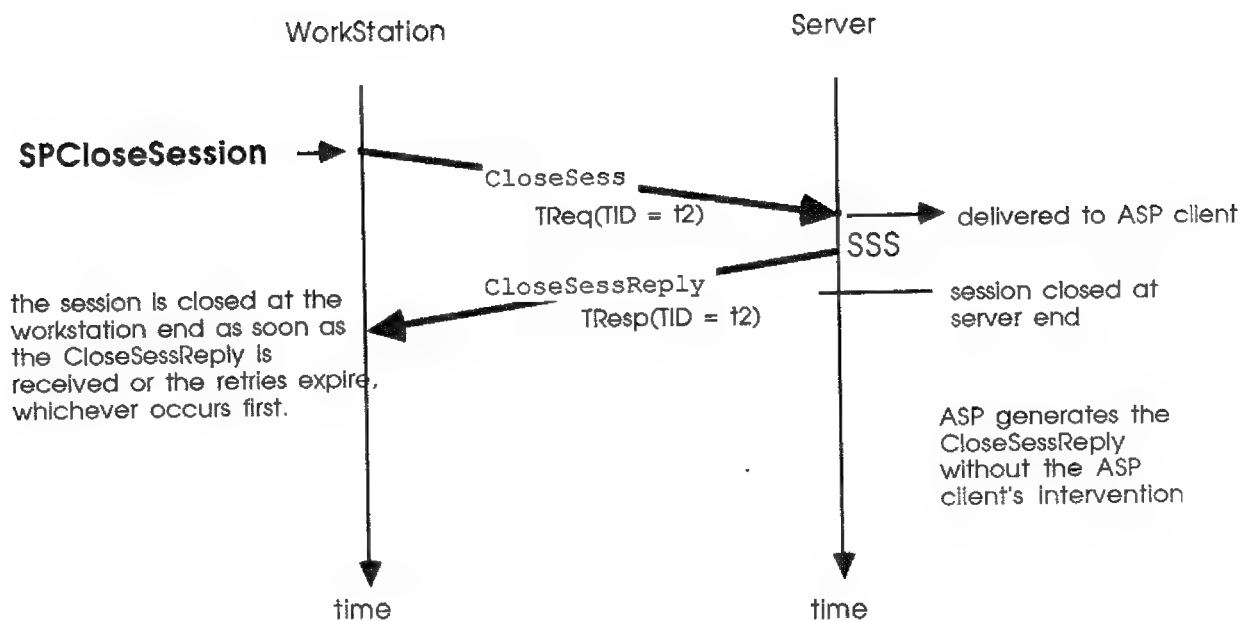


Figure XI-6. Session closing transaction (workstation initiated)

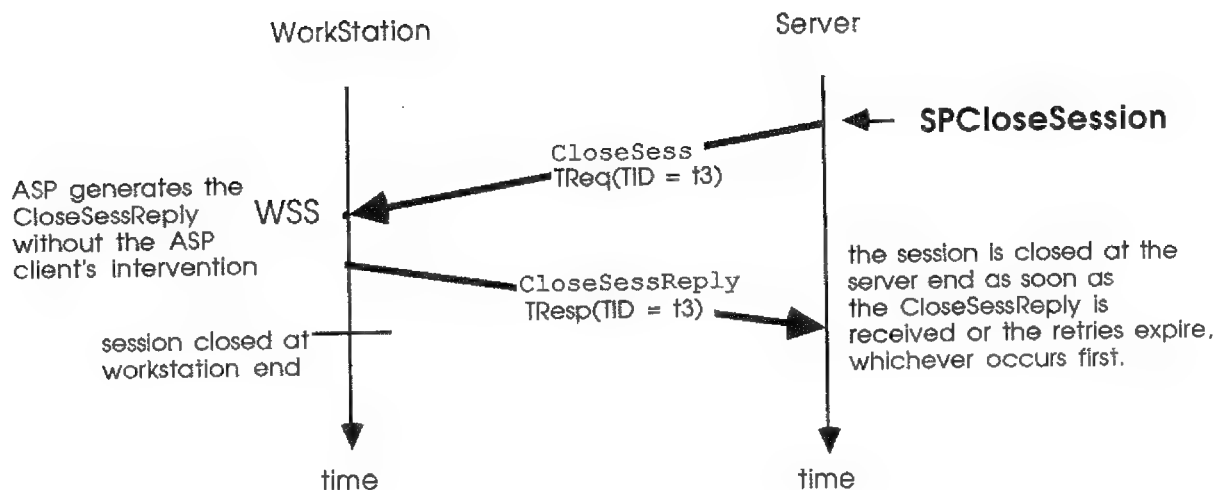


Figure XI-7. Session closing transaction (server initiated)

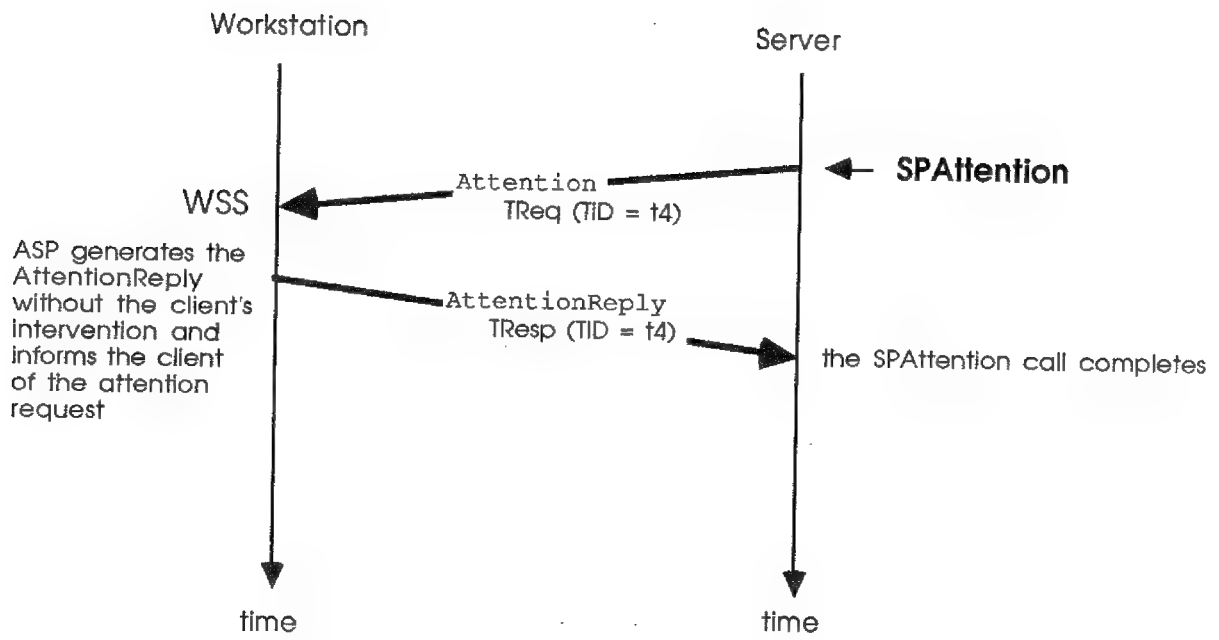


Figure XI-8. Attention request

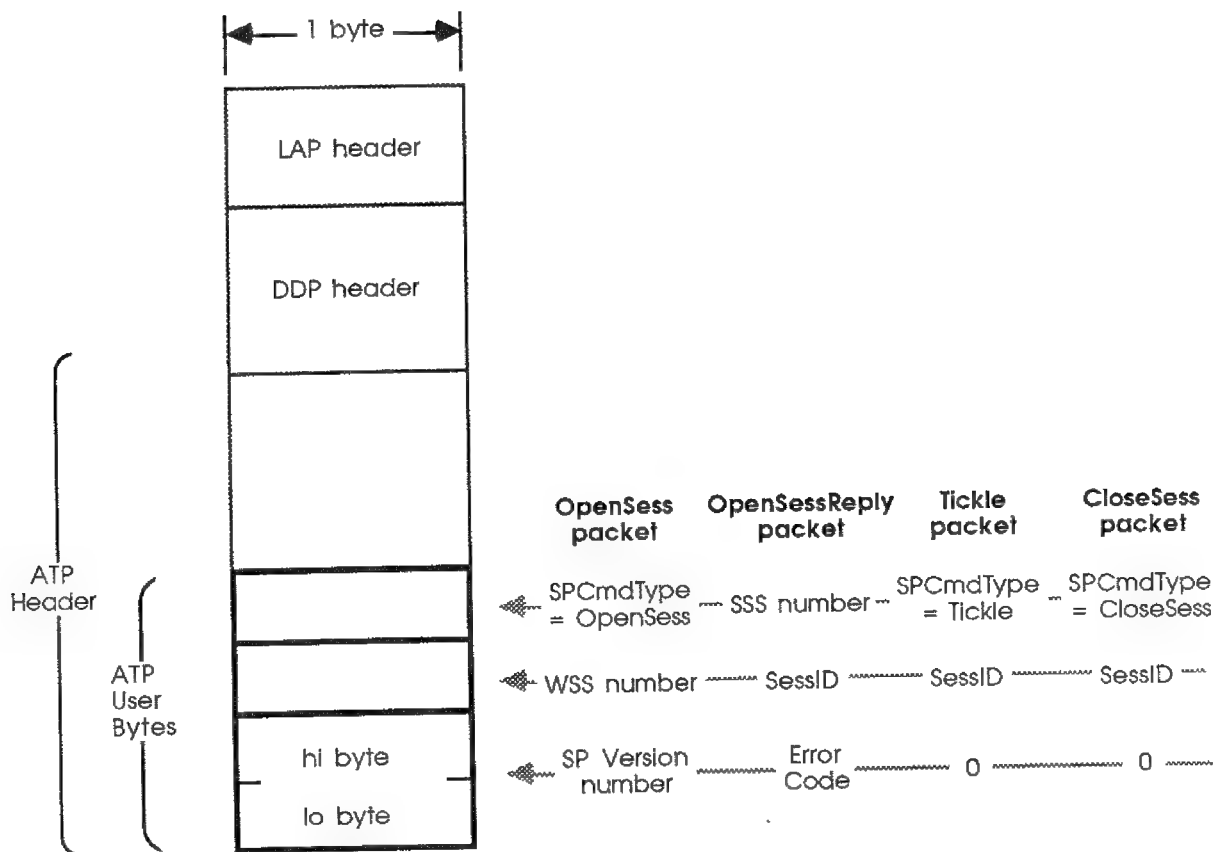
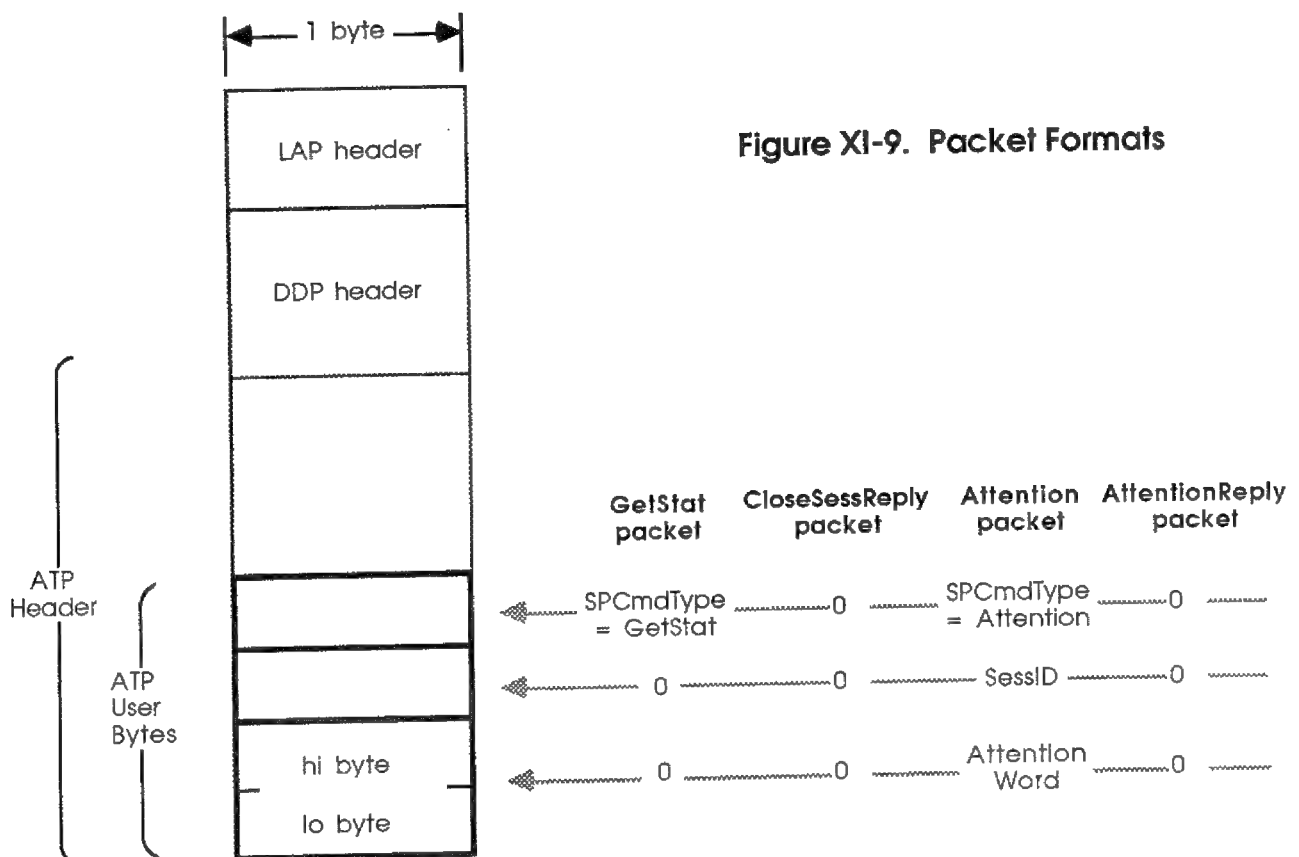


Figure XI-9. Packet Formats



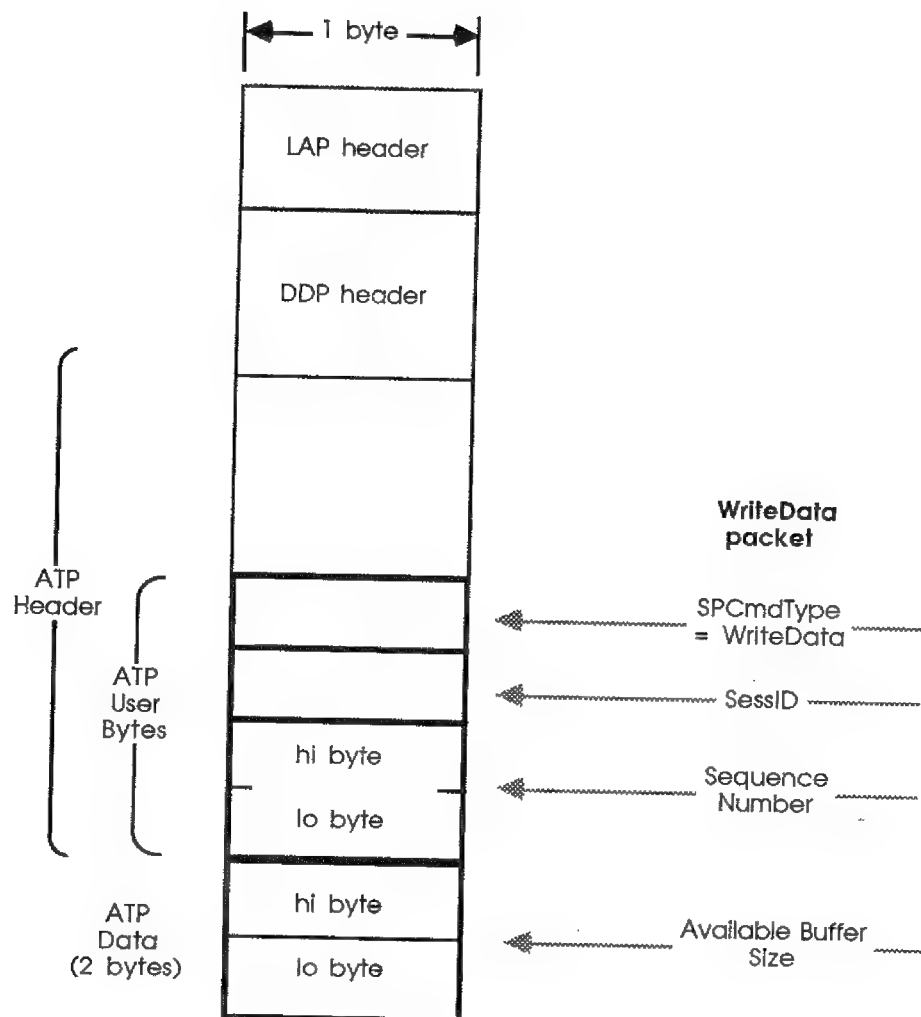


Figure XI-11. ASP Packet Formats (continued)




Figure XI-10. Packet Formats (continued)

Appendix A

AppleTalk Electrical/Mechanical Specifications

This appendix reproduces the two relevant Apple Computer specification documents for AppleTalk's electrical and physical aspects.

The first document (Apple document number 062-0190-B) is titled "AppleBus Electrical/Mechanical Specification." The use of the name "AppleBus" is historical and should be read in this context as "AppleTalk." This document provides the detailed electrical specifications as well as cable and connector characteristics.

The second document (Apple document number 157-0049-C) is titled "Transformer Specification." It provides the detailed specification of the transformer used in the AppleTalk connection module.



These documents are reproduced here for the convenience of developers wishing to implement AppleTalk compatible connections on their devices.

REV.	ZONE	ECO #	REVISION	APPD	DATE
B		J187	REDRAWN WITH CHANGES		
C	pg. 1 pg. 2 pg. 3 pg. 4 pg. 5	J386	"APPLETALK" WAS "APPLEBUS" LEAKAGE IND. WAS 15uH MAX CAPACITANCE WAS 5pF MAX ADDED "TOP VIEW" TO BUILD DETAIL ADDED DETAILS AND DIMENSIONS, CORRECTED DIM. AND LINE LOCATIONS	J2C	1/85

Transformer Specification

157-0049-C

1/5

		METRIC		 apple computer inc.	
<small>DIMENSIONS ARE IN MILLIMETERS TOLERANCES</small> X _____ X.X _____ <small>ANGLES</small> <small>TRUNCATED DRAWING</small>				<small>NOTICE OF PROPRIETARY PROPERTY</small> THE INFORMATION CONTAINED HEREIN IS THE PROPRIETARY PROPERTY OF APPLE COMPUTER, INC. THE POSSESSOR AGREES TO THE FOLLOWING: 1. TO MAINTAIN THIS DOCUMENT IN CONFIDENCE 2. NOT TO REPRODUCE OR COPY IT 3. NOT TO REVEAL OR PUBLISH IT IN WHOLE OR PART	
<small>MATERIAL</small> GARDNER		<small>FINISH</small> 5G		<small>DATE</small> 8/84	
<small>DESIGN</small> J2C		<small>CHKD BY</small> J2C		<small>DATE</small> 8/84	
<small>APPROVED</small> G. Crow		<small>DATE</small> 1/85		<small>TITLE</small> Transformer	
<small>DESIGNED BY</small> G. Crow		<small>DATE</small> _____		<small>PART NO.</small> 157-0049-C	

5

1.0 Description

This transformer is used in the Appletalk connection module to give isolation between the Appletalk cable and the devices which are connected to the cable.

The transformer is a 1:1 turns ratio transformer with tight coupling between primary and secondary, and electrostatic shielding to give excellent common mode isolation.

The primary is wound as two windings of #32 AWG wire in series with one wound below the secondary and one above it. The secondary is a single continuous winding of #32 wire.

2.0 Environmental

The transformer shall operate properly and meet its specifications under the following environmental conditions:

Operating Temperature:	0 to 70 ° C
Storage Temperature:	-40 to 70 ° C
Relative Humidity:	5 to 95 %
Altitude:	0 to 4572 meters

In addition, the transformer must meet the Apple Computer shock and vibration requirements while mounted on a printed circuit board and tested to Apple specification number 062-0086.

3.0 Mechanical Strength and Workmanship

The transformer winding assembly, pins, mounting plate, core, and clamp shall be securely mounted and rigid with respect to each other.

The pins must be easily solderable; solderability must meet EIA-RS-186-9E. All components shall be free of undue mechanical stresses.

4.0 Identification Markings

The transformer shall be marked with the manufacturer's name or identification number, date of manufacture, country of origin, manufacturer's part number, and the Apple part number. The markings shall be clearly legible after typical assembly, wave-soldering, and cleaning processes, and after exposure to the above mentioned environmental extremes. The markings may be placed in any convenient and visible location on the transformer.



apple computer inc.

SIZE
A

DRAWING NUMBER
157-0049-C,

SCALE:

SHEET 2 OF 5

5.0 Electrical Specifications

Core Material:	Siemens B65651-K000-R030 (or equivalent)
Bobbin:	Siemens B65652-PC1,L (or equivalent)
Retaining Clip:	Siemens B65653-T (or equivalent) with washer if required
Magnetizing Inductance:	15 mH minimum @ 10 KHz, 1 V RMS measured between pins 1-3
Leakage Inductance:	17 μ H max @ 10 KHz (measured between pins 1-3 with pins 5-6 shorted.)
Resistance:	1.5 Ω max (measured between pins 1-3, and 5-6)
Capacitance:	6 pF max (primary to secondary with electrostatic shield and core guarded)
Turns Ratio:	1:1 accurate to the nearest 1/2 turn
Hi Pot:	From W2 to core, shield and primary 1000 VDC for one minute with no significant leakage current.
Magnitude of Impedance (IZI)	10K Ω minimum @ 250 KHz Measured with HP 4800A Vector Impedance Meter between pins 5-6

Notes: Different permeability core material may be used, and turn count may be modified as long as all electrical specifications are met.



apple computer inc.

SIZE
A

DRAWING NUMBER
157-0049-C

SCALE:

SHEET 3 OF 5

5.0 Interconnection

5.1 AppleBus devices are connected to the AppleBus by a connection module which contains a transformer, DB-9 connector at the end of an 450 millimeter cable, and two 3-pin miniature DIN connectors as shown in Figure 5.1. Each 3-pin connector has a coupled switch. If both connectors are used, the switches are open, but if one of the connectors is not used, a 100 ohm termination resistor (R2) is connected across the line. The use of the connection module allows the AppleBus device to be removed from the system by disconnecting it from the module without disturbing the operation of the bus. R3 and R4 increase the noise immunity of the receivers, while R5 and C1 isolate the frame grounds of the AppleBus devices and prevent ground loop currents. The resistor (R1) provides static drain for the cable shield to ground.

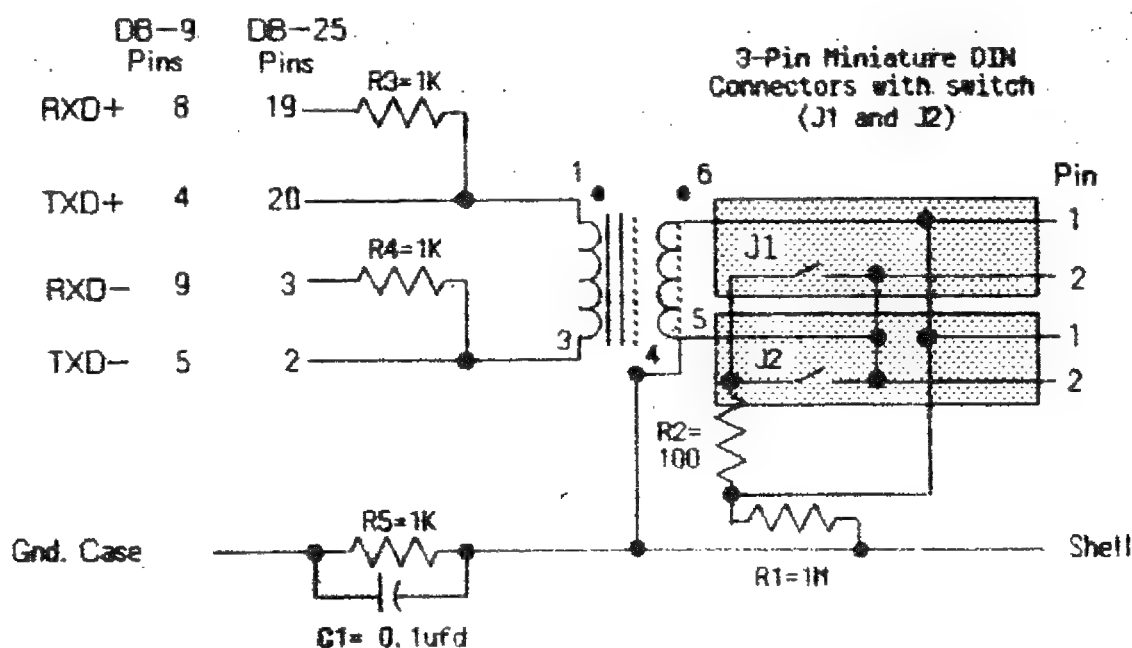


Figure 5.1 AppleBus Connection Module



apple computer inc.

SIZE
A

DRAWING NUMBER

062-0190-B

SCALE: _____

SHEET 4 OF 10




FIGURE 2.0 MECHANICAL SPECIFICATION



apple computer inc.

SIZE
A

DRAWING NUMBER
157-0049-C

SCALE: 2X



SHEET 5 OF 5

REV.	ZONE	ECO #	REVISION	APPD	DATE
B		J186	REDRAWN WITH CHANGES	<i>LLC</i>	8/84

AppleBus Electrical/Mechanical Specification

062-0190-B

1/10

		METRIC		 apple computer inc.	
<small>DIMENSIONS ARE IN MILLIMETERS TOLERANCES</small> X ± _____ ANGLES _____ XX ± _____ ANGLES _____ <small>(DO NOT SCALE DRAWING)</small>				NOTICE OF PROPRIETARY PROPERTY THE INFORMATION CONTAINED HEREIN IS THE PROPRIETARY PROPERTY OF APPLE COMPUTER, INC. THE POSSESSOR AGREES TO THE FOLLOWING: 1. TO MAINTAIN THIS DOCUMENT IN CONFIDENCE 2. NOT TO REPRODUCE OR COPY IT 3. NOT TO REVEAL OR PUBLISH IT IN WHOLE OR PART	
MATERIAL _____		FINISH _____		TITLE	
DRAWN GARDNER	8/84	CHECKED <i>LLC</i>	8/84	SPECIFICATION, APPLEBUS ELECTRICAL/MECHANICAL	
ENG. APPROV. <i>LLC</i>	8/84	DESIGNED <i>LLC</i>	8/84		
RELEASE <i>LLC</i>	8/84	ORIG. DIS. M. A. C.			
DESIGNER _____	SCALE _____	SIZE A	DRAWING NUMBER 062-0190-B		SHT 1/10

1.0 Introduction

AppleBus is a serial interconnection system for all Apple Computers and several future peripheral products. It provides a common, convenient, inexpensive, expansion capability for computer products. In summary, AppleBus is a multi-drop, balanced, transformer isolated, serial communication system designed to connect 32 devices at 230.4 Kbaud over a total distance of up to 300 meters.

2.0 References

AppleBus Link Access Protocol, specification number 062-0214.

EIA Standard RS-422, Electrical Characteristics of Balanced Voltage Digital Interface Circuits.

Fairchild Semiconductor "Interface, Line Drivers and Receivers", 1975.

3.0 Summary of Features and Performance

- * 32 total devices
- * Shielded, twisted pair, connectorized interconnection; easy user configuration, maximum total cable length of 300 meters
- * 230.4 Kbaud communication, SDLC frame format, FMO modulation
- * Balanced signalling using standard RS-422 driver (26LS30) and receiver (26LS32) I.C.'s
- * Transformer isolation for excellent noise and static discharge immunity
- * Self-configuring, no user switches or action to identify devices
- * Passive drops, devices may be disconnected, and one may fail without disturbing communication



apple computer inc.

SIZE
A

DRAWING NUMBER
062-0190-3

SCALE: _____

SHEET 2 OF 10

4.0 Signalling

- 4.1 AppleBus devices send and receive data over a single pair of wires connected to each device. Two connectors on each device allow the user to easily connect devices together by means of a simple cable (see Section 5.0 Interconnection). Balanced, transformer coupled signalling is used to reduce both RFI and noise susceptibility.
- 4.2 Each device has a single driver and receiver. The receivers are always connected to the system and pass all AppleBus data to the controller. Only one driver at a time is enabled. Software controls which device may transmit data (see Protocol Specification). The driver and receiver descriptions and specifications are given in the electrical section, 6.0.
- 4.3 The signal on AppleBus is encoded with FMD (bi-phase space). This insures that clock information can be recovered at the receivers. In FMD, a transition occurs at the beginning of every bit cell. A "0" is represented by an additional transition at the center of the bit cell, and a "1" is represented by no transition at the center of the cell.




Figure 4.1 FMD Coding (230.4 K baud \pm 1%)

- 4.4 Synchronization time for the clock recovery system is provided by the transmission of 14 consecutive 1 bits directly preceding the data frame. Frame format is covered in the Protocol Specification.

5.0 Interconnection

5.1 AppleBus devices are connected to the AppleBus by a connection module which contains a transformer, DB-9 connector at the end of an 450 millimeter cable, and two 3-pin miniature DIN connectors as shown in Figure 5.1. Each 3-pin connector has a coupled switch. If both connectors are used, the switches are open, but if one of the connectors is not used, a 100 ohm termination resistor (R2) is connected across the line. The use of the connection module allows the AppleBus device to be removed from the system by disconnecting it from the module without disturbing the operation of the bus. R3 and R4 increase the noise immunity of the receivers, while R5 and C1 isolate the frame grounds of the AppleBus devices and prevent ground loop currents. The resistor (R1) provides static drain for the cable shield to ground.

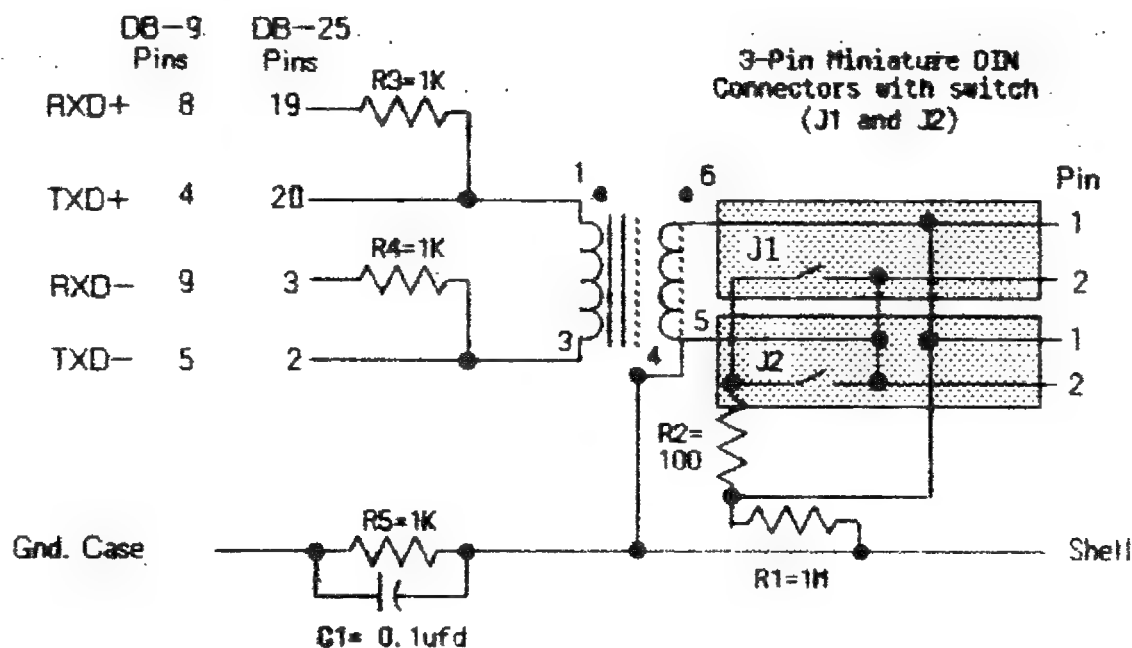


Figure 5.1 AppleBus Connection Module



apple computer inc.

SIZE
A

DRAWING NUMBER
062-0190-B

SCALE: _____

SHEET 4 OF 10

- 5.2 Individual AppleBus devices are connected together by a twisted, shielded pair cable. The cable is available in standard lengths with connectors on each end, or in 150 meter bulk rolls. The maximum length of cable for the bus is limited to a total of 300 meters.

Cable Specification

Conductors:	22 AWG stranded 17 ohm per 300 meters
Shield:	85% coverage braid
Impedance:	78 ohm
Capacitance:	68 pF per meter
Rise Time:	175 ns 0 to 50% at 300 meters
Diameter:	4.7 mm (0.185 inches) maximum



apple computer inc.

SIZE
A

DRAWING NUMBER
062-0190-B

SCALE: _____

SHEET 5 OF 10

5.3 The AppleBus connector is a miniature 3 pin circular connector similar to Hosiden connector number TCP8030-01-010.

Pin 1 AppleBus - Plus
Pin 2 AppleBus - Minus
Pin 3 Unused
Shell Shield




Figure 5.2 Connector Pin Assignment (looking into connector)

5.4 The interconnecting cable is wired "one-to-one" as shown below.




Figure 5.3 Interconnecting Cable Connection



apple computer inc.

SIZE
A

DRAWING NUMBER

062-0190-B

SCALE:

SHEET 6 OF 10

6.0 Electrical Specification

- 6.1 The recommended driver is the 26LS30 used with both +5V and -5V as power supplies, and the mode control connected to give differential outputs (Mode voltage low).

The outputs of the driver are coupled to the connector through a "degitch network" which consists of a T-network of two 20 to 30 ohm resistors and a 150 to 300 picofarad capacitor. Use of the network gives two major advantages. The first is that the high frequency components of the signal are attenuated both going onto and off of the cable thus reducing RFI and also static susceptibility. The second advantage is that at least one driver on the network can fail without causing the network to fail (as long as it fails in one state and doesn't broadcast trash).

Those who wish to use standard RS-422 specified components (power supplies of +5 volts and ground) may do so if the degitch circuits are not used, but should be aware that the driver must drive a cable impedance of 39 ohms (middle of a 78 ohm cable).

- 6.2 The receiver is the 26LS32. Both inputs of the receiver are connected through degitch networks.

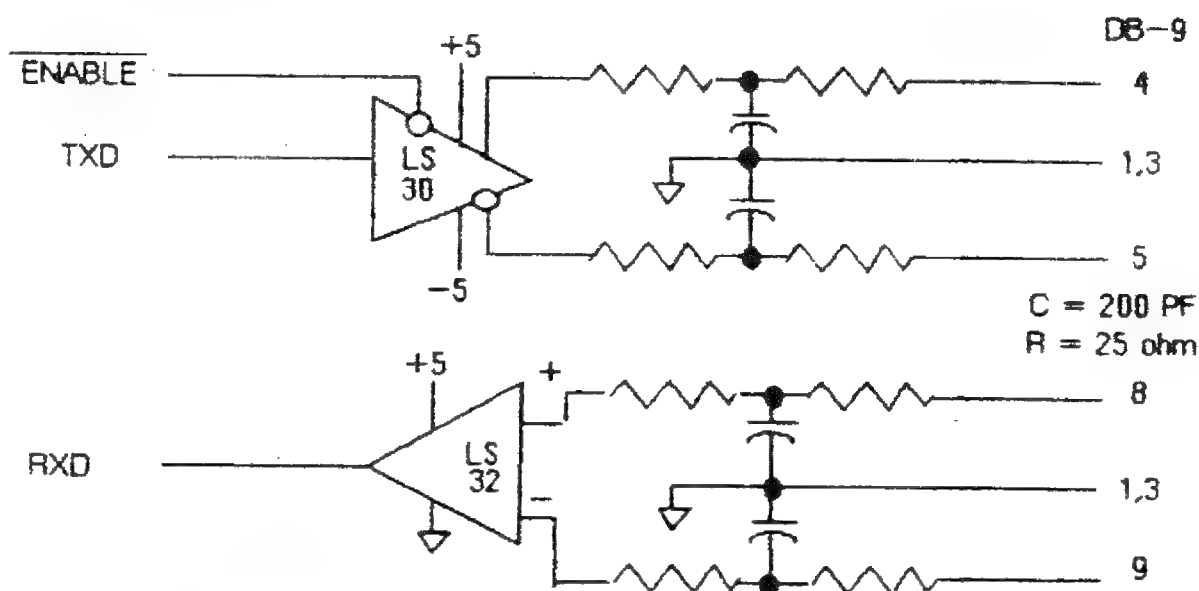


Figure 6.1 Driver and Receiver Connection



apple computer inc.

SIZE
A

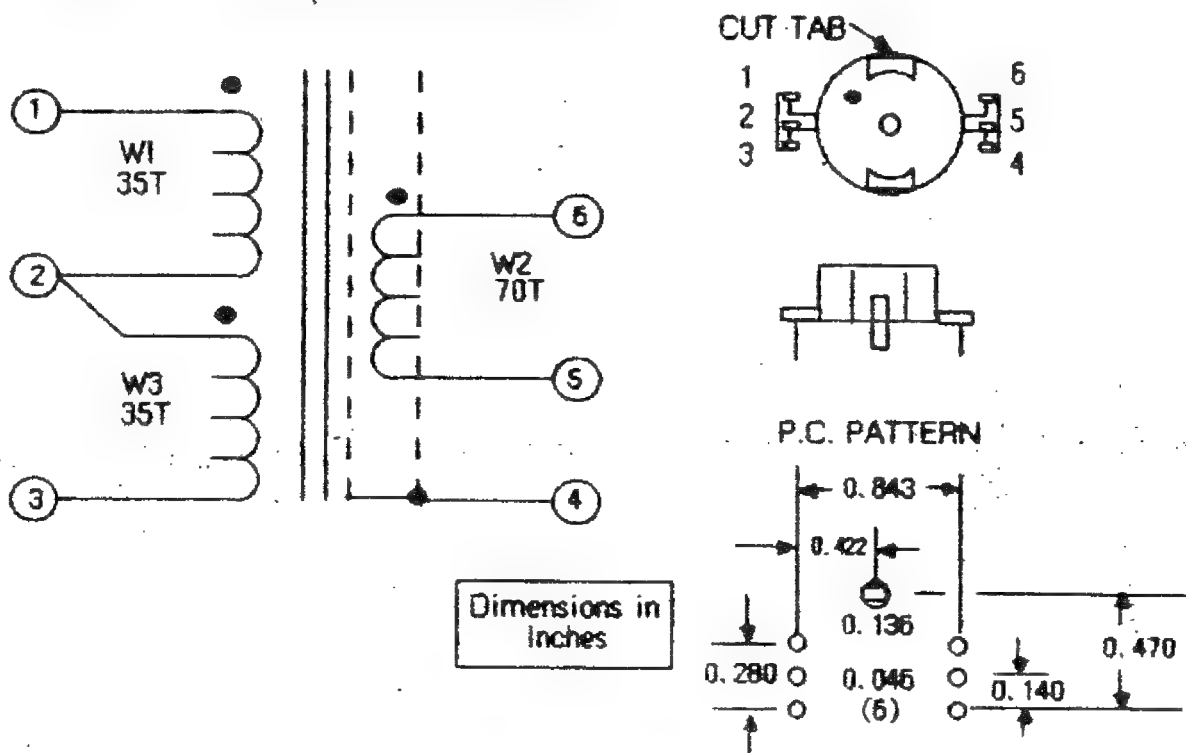
DRAWING NUMBER
062-0190-B

SCALE: —

SHEET 7 OF 10

6.3 The transformer is a 1:1 turns ratio transformer with tight coupling between primary and secondary, and electrostatic shielding to give excellent common mode isolation.

The primary is wound as two windings of #32 AWG wire in series with one wound below the secondary and one above it.



Specifications:

Core Material:	Siemens B65651-K000-R030 (or equivalent)
Bobbin:	Siemens B65652-PC1,L (or equivalent)
Retaining Clip:	Siemens B65653-T (or equivalent)
Magnetizing Inductance:	20 mH minimum
Leakage Inductance:	15 μ H max
Capacitance:	5 pF max (primary to secondary with electrostatic shield and core guarded)

Figure 6.2 Transformer Specification



apple computer inc.

SIZE
A

DRAWING NUMBER
062-0190-B

SCALE: _____

SHEET 8 OF 10

- 6.4 Each end of the cable must be terminated. The AppleBue connection module is designed such that a 100 ohm resistor is connected across the line if one of the two connectors is not used. A 100 ohm resistor is used even though the characteristic impedance of the line is 78 ohms because it gives adequate termination and minimizes resistive losses.
- 6.5 The cable shield should be grounded to Earth ground (frame ground) at each AppleBus device. This ground is necessary to prevent excessive RFI. A resistor and capacitor are included in each connection module to isolate the cable shield from the ground connection at 60 Hz while offering a low impedance connection to high frequency noise. This connection scheme allows ground connection for RFI purposes without risking high currents flowing in the cable due to differing ground potentials. In the U.S., the green wire available in most outlets is suitable for frame grounding.



apple computer inc.

SIZE
A

DRAWING NUMBER
062-0190-B

SCALE: _____

SHEET 9 OF 10

7.0 Documentation

The following is a list of all AppleBus hardware drawings.

Drawing Description	Drawing Number
AppleBus Connection Module Cover (Top)	815-0839
AppleBus Connection Module Cover (Bottom)	815-0838
AppleBus Cable Extender (Coupler)	519-0300
AppleBus Cable with Molded Plugs	590-025X
DB-9 Connection Cable Assembly	590-0254
DB-25 Connection Cable Assembly	590-0253
AppleBus Board Film Artwork	820-0135
AppleBus Connection Module FCC Label	825-1008



apple computer inc.

SIZE
A

DRAWING NUMBER
062-0190-B

SCALE: ———

SHEET 10 OF 10

Appendix B

Miscellaneous AppleTalk Parameters

This appendix summarizes various numerical quantities used in the AppleTalk protocols. It provides a single reference source for this information. This information is organized into subsections, one for each relevant protocol.

The notation used in this section is: \$.... represents hexadecimal, %.... represents binary, all other numbers are decimal.

ALAP Parameters

ALAP protocol type field values:

\$00 -	invalid ALAP protocol type value (not to be used)
\$01 through \$7F -	valid ALAP protocol type values for use in ALAP client packets;
\$01 through \$0F -	reserved for Apple Computer's use only
\$01 -	DDP short header packet
\$02 -	DDP long header packet
\$0F -	experimental ALAP packet (used by Apple only)
\$80 through \$FF -	reserved for ALAP control frames
\$81 -	ALAP ENQ packet
\$82 -	ALAP ACK packet
\$84 -	ALAP RTS packet
\$85 -	ALAP CTS packet

This information is summarized pictorially in Figure B-1.

Timing constants used by ALAP:

There are three constants used by ALAP:

Inter-Frame Gap -	less than 200 microseconds
Inter-Dialogue Gap -	at least 400 microseconds
IDG slot -	100 microseconds

ALAP frame parameters:

Flag byte used for framing an ALAP packet = %01111110
Number of flag bytes needed at the start of frame = two or more
Number of bits in abort sequence - 12 to 18
Maximum number of data bytes in ALAP frame (not including headers, etc.) = 600

DDP Parameters

DDP packet parameters:

ALAP protocol type value for short header DDP packet = 1
ALAP protocol type value for long header DDP packet = 2
Maximum number of data bytes in a DDP packet = 586 bytes

DDP protocol type field values:

\$00 -	invalid DDP protocol type value (not to be used)
\$01 through \$FF -	valid DDP protocol type values for use in DDP client packets;
\$01 through \$0F -	reserved for Apple Computer's use only
\$01 -	RTMP response or data packet
\$02 -	NBP packet
\$03 -	ATP packet
\$04 -	EP packet
\$05 -	RTMP request packet
\$06 -	ZIP packet
\$07 -	ADSP packet

This information is summarized pictorially in Figure B-2.

Socket numbers:

\$00 -	invalid (do not use)
\$FF -	invalid (do not use)
\$01 through \$FE -	valid DDP sockets
\$01 through \$7F -	statically assigned sockets
\$01 through \$3F -	reserved for Apple Computer's use only
\$40 through \$7F -	experimental use only (<u>not to be used in released products</u>)
\$01 -	RTMP socket
\$02 -	Names Information Socket
\$04 -	Echoer Socket
\$06 -	Zone Information Socket

This information is summarized pictorially in Figure B-3.

RTMP parameters

RTMP Socket = socket number 1
Maximum number of hops = 16

RTMP packet parameters:

DDP protocol type value RTMP response or data packets = 1
DDP protocol type value RTMP request packets = 5
RTMP request packet command field values:
1 = Request

RTMP Timers:

Send-RTMP Timer = 10 seconds

Validity Timer = 20 seconds

Timer for aging A-Bridge in a non-bridge node = 90 seconds

NBP parameters

Names Information Socket = socket number 2

Maximum number of characters in object, type or zone fields = 32

Wildcards in NBP names:

- * used only in the zone field of an entity name to mean the zone of the packet's sender,
- = used as the object and/or type field of an entity name to mean all objects and/or all types (can not be used as the zone field)

NBP packet parameters:

DDP protocol type value for NBP packets = 2

NBP Control field values:

1 = BrRq

2 = LkUp

3 = LkUp-Reply

ATP parameters

ATP packet parameters:

DDP protocol type value for ATP packets = 3

Function code values:

%01 = TReq

%10 = TResp

%11 = TRel

Maximum size of data in an ATP packet = 578 bytes

ATP Timers:

Release timer = 30 seconds

ZIP parameters

Zone Information Socket = socket number 6

ZIP packet parameters:

DDP protocol type value for ZIP packets = 6 (except GetMyZone and GetZoneList packets)

ZIP Command values:

Inside AppleTalk - July 86

- 1 = ZIP query
- 2 = ZIP reply
- 3 = ZIP takedown
- 4 = ZIP bringup
- 7 = ZIP GetMyZone
- 8 = ZIP GetZoneList

ZIP Timers:

Query retransmission timer = 10 seconds

ZIP bringback timer = 10 minutes

PAP parameters

PAP packet parameters:

Maximum data size in packet = 512 bytes

PAP Type values:

- 1 = OpenConn
- 2 = OpenConnReply
- 3 = SendData
- 4 = Data
- 5 = Tickle
- 6 = CloseConn
- 7 = CloseConnReply
- 8 = SendStatus
- 9 = StatusReply

Maximum length of status string = 255 bytes (not including the length byte)

PAP Timers and Retry Counts:

OpenConn request ATP retry timer = 2 seconds

OpenConn request retry count = 5

Tickle timer = 60 seconds

Connection timer = 2 minutes

SendData request retry timer = 15 seconds

EP parameters

Echoer Socket = socket number 4

EP packet parameters:

DDP protocol type value for EP packets = 4

Echo Function values:

- 1 = Echo request
- 2 = Echo reply

Maximum data size = 585 bytes

ASP parameters

ASP packet parameters:

SPCmdType values:

- 1 = CloseSession
- 2 = Command
- 3 = GetStat
- 4 = OpenSess
- 5 = Tickle
- 6 = Write
- 7 = WriteData
- 8 = Attention

ASP Timers:

Tickle timer = 30 seconds
Session maintenance timer = 2 minutes

Appendix C

AppleTalk Peek

The *AppleTalk Peek* program is a network tool used to monitor packet traffic on an AppleTalk network. Peek runs on any Macintosh except the Macintosh XL *with AppleTalk connected via the Printer port*, and can record all packets seen on the bus. In addition, it can detect certain errors, measure packet arrival times, and display packet data in hexadecimal and ASCII format.

Peek has enough queue space to hold a large number of packets. This queue is used in a circular fashion, so that Peek can continue to monitor packets even after the queue has been filled. Older packets are discarded to make room for newer ones.

This document describes version 3.0 of Peek.


The Window

When Peek is started, the program's window is drawn. It contains the control buttons, menus, and information display areas described below.



Peek is always in one of two states: recording or displaying packets. When the program is first started, it is in the record state. The **START** and **STOP** buttons are used to initiate and terminate a recording session, during which Peek listens on the bus and records traffic.

When the **START** button is pressed, packet recording is enabled. The button becomes gray to indicate that Peek is recording (see Figure 1). Peek's internal buffers are cleared, and packets from a previous session are lost. The **STOP** button halts recording and causes packets to be displayed, if any were recorded during the session (see Figure 2).



When the **STOP** button is pressed, the "Pkts in Q" box shows the number of packets in Peek's queue. This box is not dynamically updated during a recording session, since queue wraparound makes this determination difficult. The word "Sampling" appears here while recording, as an indication that Peek is monitoring the bus.

The size of the queue is determined by the amount of free memory, so that Peek running on a 512K Macintosh will be able to record more packets than a 128K Macintosh. There is, however, a limitation of 190,000 bytes total in the receive buffer.

Pkts Rcvd
0

The total number of packets seen by Peek since the start of the recording session appears in the "Pkts Rcvd" box. This count is updated dynamically, and hence provides a rudimentary visual indication of bus traffic. Since Peek's queue can wrap around and "forget" old packets, this count may be greater than the number of packets stored in the queue at that time.

CRC errors:	0
Overruns:	0
Time Outs:	0

This box displays the tally of errors during a recording session. Peek can detect three types of errors:

CRC errors are noted when the 16-bit Frame Check Sequence (FCS) at the end of the ALAP frame does not match the calculated FCS. This indicates a possible error in transmission (due to noise on the bus, collisions, etc.).

Overflow errors occur when the receiver reads bytes too slowly from the Macintosh's Serial Communications Controller (SCC), and this chip's three-byte FIFO buffer overflows. Note that if the node running Peek detects an overrun error, it does not necessarily mean that this will be the case for other nodes. This error is sometimes detected as a by-product of collisions on the bus.

Timeout errors are flagged when a byte is expected on the bus (the end of the packet has not been seen, yet a byte does not appear on the bus within about 400 microseconds of the previous byte). Every byte received thus far will be stored and displayed as "the packet", even though the true packet end was not detected. This usually indicates a problem in the packet sender's hardware, but it may also be a by-product of collisions on the bus.

The error fields are updated "on the fly" as packets are recorded, and are a measure of the total number of errors seen by Peek. Therefore, in a long recording session, it is possible, due to queue wraparound, for a packet to be received in error and not appear in the packet display, although the error is noted in the box.

Rcv Status
Receiving LAP packets.
Receiving RTMP packets.

The "Rcv Status" box indicates to the user the status of two different types of packet reception filters. If a status line is displayed, then that filter is disabled. The default is set to 'filter enabled'.

Display box

This box is used to view packets which were saved during the recording session. Each packet is preceded by a banner line (see Figure 2) in boldface which includes, from left to right:


- a set of square brackets which may contain blanks or one of the following characters: **B** if the packet was a broadcast, **C**, **O**, or **T** if a CRC, overrun, or timeout error, respectively, was detected for that packet;
- the source **S** and destination **D** node IDs of the packet, in decimal format;
- the packet's arrival time **T** in milliseconds (measured relative to the first packet stored in the queue);
- the time since the previous packet's arrival (delta time or ΔT);
- the packet's sequence number in parentheses (in the order in which they were received, starting with zero for the oldest packet in the queue);
- the calculated length **L** of the packet (number of bytes in decimal, not including the FCS).

Following this banner line are two displays of the packet's contents, in hexadecimal on the left and the corresponding ASCII (if printable) on the right. A period is substituted for any unprintable character. Note that the FCS bytes are not shown.

On the right side of the display box is a scroll bar which is enabled whenever there is more to be displayed than will fit in the box. The user can scroll through the display of packets by using the scroll bar's up and down arrows, or by dragging the thumb. Clicking in the gray area above or below the thumb shifts the display backwards or forwards one page at a time. As an alternative, holding down the option key while clicking in the grey area will scroll one entire packet at a time. This is useful for scrolling past large packets.

Menus



Menu:



The Apple menu, as usual, is used to invoke a variety of desk accessories. Choosing **About Peek...** will cause Peek to display some descriptive information, including its version number and the size of the queue in bytes.

File

Menu:



The second menu is the **File** menu, as seen above. Items in this menu allow you to save packet images, print them out, or just exit the application.

Before packets are saved or printed, you may use the **Set Selection Range...** menu item. This command allows you to select a range of packets in the queue to be later saved or printed. The default selection is set to "all" when Peek starts, however you may change this to any range valid for the current buffer. Please note that this selection range is in effect until it is changed or the program restarts.

Selecting **Save Selection** will cause Peek to bring up the standard Save dialog, in which case you may save away a copy of selected packets into a file on disk. If you have used the **Set Selection Range...** command to make a selection range, only those packets will be saved onto your disk.

If there is not enough room on the disk to save all the packets, Peek will write as many as will fit, and notify the user that a "Write Error: -34" occurred. This means that the disk is full; you may wish to put a different disk in the drive and try again. The file containing saved packets will be of type "TEXT" and creator "EDIT".

By choosing **Print Selection**, the selected packet range will be printed onto your printer. As in **Save Selection**, if you have chosen a selection range, only those will be printed. The printer that will be used is the one that was selected using the Choose Printer or Chooser desk accessories.

Note that the **Short Format** option (described later) has no effect on packets written to a file or to the printer; the long format is always used.

Edit Menu:

Apple	File	Edit	Search	Control
		Cut	⌘H	
		Copy	⌘C	
		Paste	⌘U	

The **Edit** menu is used only in conjunction with the **Find Pattern** feature in the **Search** menu, described below. It allows you to use the standard text edit commands to create a string for which to search.

Search Menu:

Apple	File	Edit	Search	Control
			Find Pattern	⌘F
			Find Same	⌘S
			Find Overrun	
			Find CRC Error	
			Find Timeout	


The **Search** menu is used to look for a particular hexadecimal or ASCII string within the recorded packets. Selecting Find Pattern will cause the Find window to appear. The standard text editing features available in the Edit menu may be used. Hex strings must be a sequence of bytes, each specified as a dollar sign (\$) followed by a two-digit hex number. A wild-card may be specified by the command key-equals sign combination. This will appear in the Find window as "Ω", and will match one or more characters of any value.

When the string has been entered, hit Return or the Find Next button. Peek will begin searching from the first packet appearing in the display box. The display will be scrolled down to the packet containing the string, if found, and the string will be highlighted. Otherwise, Peek will inform you that the string was not found. Selecting Find Same (or the equivalent command key-S combination) will cause Peek to look for the next occurrence of the same string, starting from the current packet. Note that searches are always made case-sensitive.

Find Overrun, **Find CRC Error**, and **Find Timeout** work in a similar fashion. Selecting one causes Peek to search, starting from the first packet in the display box, for a packet exhibiting the particular error. If found, the display is scrolled to bring that packet to the top of the box (unless it is too close to the last packet to scroll up to the top). Since the error counts are cumulative from the time the START button was pressed, packets with

errors may not always appear in the queue (if they were discarded to make room for newer packets). If there are no errors of that type, the menu item is dimmed.

Control Menu:



The **Control** menu allows you to manipulate packet information. When **Display Short Format Pkts** is selected, a check mark appears next to the menu item and packets are displayed in a more compact form. Only the banner line and the first line (up to the first 16 bytes) of each packet will be shown. The display can be scrolled as before. Choosing the menu item again will change the display back to long format. This item is inactivated during a recording session.

When **Receive LAP Control Pkts** is selected, all ALAP control packets seen on the bus will be recorded. These are defined as any packets whose ALAP type field is \$80 through \$FF hex (most significant bit set). Such packets will be recorded only when this option is selected. If their reception is enabled in the middle of a recording session, any ALAP control packets already seen on the bus will not have been saved. Changing this option while viewing the display of a previous session will have no effect on the display, but will affect the next recording session.

The **Receive RTMP Pkts** is used to filter out RTMP type packets. RTMP packets are defined as having a ALAP type field of \$01 and a DDP type field of \$01. See the appropriate section of Inside AppleTalk for more information on the RTMP protocol.

Notes

1. The newer versions of Macsbug (1/1/85 or later, with symbols) tend to slow Peek enough that it will frequently overrun. Use older debuggers on your Peek disk or none at all.
2. When Peek starts up, it checks low memory variables to see if a RAM routine (like a debugger) is hooked into the trap dispatcher. If one is found, a warning message is displayed, allowing the user to take appropriate action.
3. A node running Peek is in listen-only mode on an AppleTalk network. Such a node does not participate in the ALAP protocol and does not even consume a node ID. In fact, it is "invisible" to other nodes.
4. Peek does not use any of the standard AppleTalk drivers (e.g. the Macintosh Protocol Package), but assumes direct control of the Macintosh's AppleTalk port. However, the port is reset when Peek terminates, so it is possible to then run other AppleTalk

software without powering down the Macintosh and powering it up again. (Note that this is not true for versions of Peek older than version 2.0).

4. Peek will not run on a Macintosh XL.
5. If a packet that is longer than 4095 bytes is received by Peek, its subsequent behavior becomes unpredictable. If Peek terminates abnormally during a recording session, there is a strong probability that a node on the network has sent a packet of illegal size (e.g. a node is stuck in its transmit loop).

Acknowledgements

This program borrows ideas from a former Lisa WorkShop application developed by Jim Nichols and Steve Butterfield; in particular, the circular use of a buffer. AppleTalk Peek is a completely new program designed to exploit the Macintosh user interface. It was written by Rich Andrews and Gursharan Sidhu. Gene Tyacke, Mark Neubieser and Paul Williams have implemented many new features.

Appendix D

AppleTalk Poke

AppleTalk Poke is a Macintosh application designed for use by AppleTalk developers. It allows the user to edit/create packets and to send them out on AppleTalk. Developers are expected to use Poke to test their protocol software/hardware implementations for AppleTalk products. Poke uses the Macintosh Protocol Package (MPP) for AppleTalk access (details of MPP are discussed elsewhere). This means that the system file of the boot disk must have AppleTalk Installed (e.g. with the **Install** tool). This has been done to the disk Poke is on.

This document describes the features and use of Poke, version 3.1. It is not intended to instruct the user on the capabilities, features, or specifications of MPP or of the various AppleTalk protocols, nor does it discuss the normal use of the Macintosh's standard editing abilities. (For information on AppleTalk protocols, see the corresponding specification/description document.)

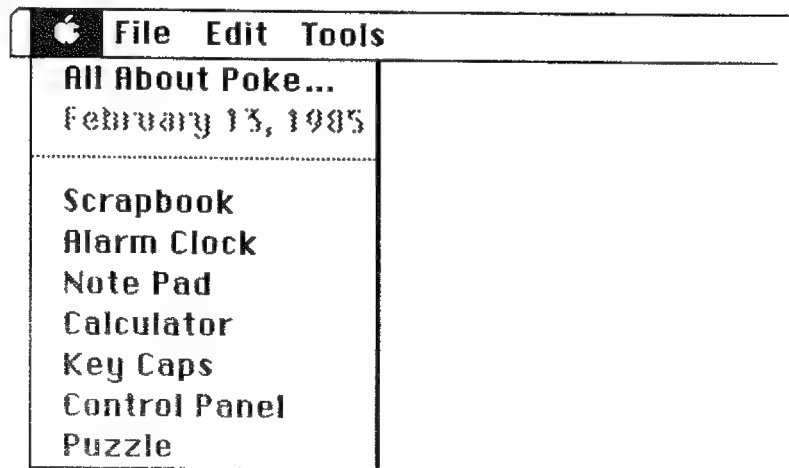
Startup

After starting Poke, the MPP driver is loaded in (if it isn't currently in memory) and the main window is brought up (Figure D-1). This window displays the Poke station's AppleTalk node ID and packet information. At this stage, the packet information indicates that no packets have been loaded into Poke (packet names are all set to empty). Next to each packet name, there is a pair of buttons labeled **EDIT** and **SEND**. Initially, all **SEND** buttons are dimmed (inactive) because no packets have been loaded. The main window includes an area for displaying any appropriate error or status messages.

The program operates in two different states. When started up, it is in the send state with the main window displayed. When any **EDIT** button is pressed, it goes into an edit state and the packet editing window is displayed (Figure D-2). In this state, the selected packet can be edited. Clicking the edit window's **OK** button will return you back to the main window and the send state.


Menus and Commands

🍏 Menu:




The 🍏 menu allows you to run an available desk accessory or to examine Poke's version information ("**All About Poke...**"). Selecting the "**All About Poke...**" command brings up an information window. Clicking the mouse or pressing a key causes this window to disappear and Poke returns to its original state.

File Menu:

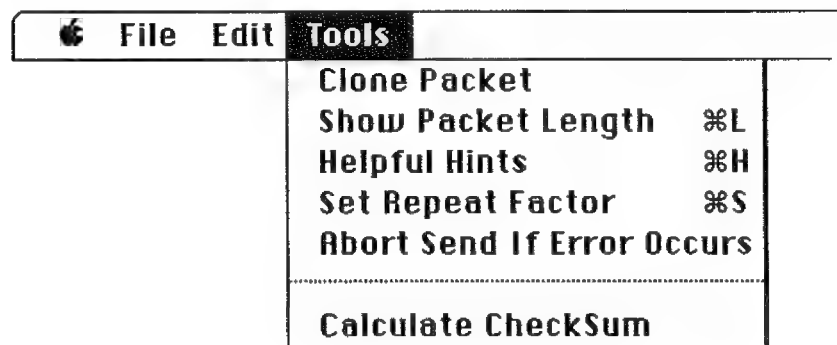


The **File** menu allows the user to load from (or save to) a file of 10 prepared or canned packets. The **Load** and **Save** operations follow the standard conventions for file loading and saving.

Note: Older versions of Poke utilize a different file format. You cannot load in packets created by those versions.

Edit Menu:

The **Edit** menu is used only while editing a packet. Please note the keyboard's optional command keys that can be used to invoke this menu's commands.

Tools Menu:

"**Clone Packet**" can only be selected from the main window. When selected, a dialog box appears which asks you for the name of the packet you wish to copy (the source). It also asks for the name of the packet which it is to be copied to (the destination). The names are searched in a top to bottom fashion starting at the top left corner of the main window (Figure D-1). The first packet whose name matches the one you entered will be chosen. If both source and destination names are found, then the source packet will be copied verbatim to the destination. Otherwise, an error message is displayed.

The "**Show Packet Length**" command can only be used while editing a packet. It returns the number of bytes in the packet's data field. This count does not include the packet's header, so the actual packet size will be larger. (See the AppleTalk protocol documentation for information on the size of the different headers.) If an error is detected while computing the length, an alert box will be displayed indicating the exact location of the error.

Note: If you have entered more data into a packet than is allowed by the corresponding protocol (LAP,DDP,ATP) then Poke will truncate the data (at the end) to the maximum allowed value.

The "**Helpful Hints**" command allows you to obtain a quick summary of editing instructions. Clicking the mouse or pressing any key will return you to the currently active Poke window.

Packets can be transmitted repeatedly at user specified intervals. The number of times a packet is transmitted and the time interval between transmissions are set by the user by selecting the "**Set Repeat Factor**" command. This command will allow you to change transmission information used by the **SEND** command (discussed later).

The delay time interval between transmissions is given in ticks (1 tick = 1/60 of a second). If you enter a number of transmissions value equal to zero, then Poke will keep sending packets out in a closed loop (i.e, indefinitely). When Poke is in such a loop, you can stop the **SEND** operation by either clicking the mouse button or by pressing a key. If you wish to send packets out at the fastest rate possible, enter a zero for the time interval. If this is done, packet statistics will not be displayed in the messages box.

Note: The user specified time interval is achieved only approximately. Network loading and ALAP overhead plus packet transmit time add to this interval.

The "**Abort Send If Error Occurs**" command is used in conjunction with the **SEND** operation. If selected, a checkmark will appear on the left side of the command informing the user that this feature is active. Now, if an error occurs while sending a packet, the **SEND** operation will abort. To deactivate this feature, select the command again and the checkmark will be removed. This command is especially useful when large numbers of packets are being sent out.

The last command, "**Calculate Checksum**", may be used in the edit window to replace the existing DDP checksum field with an updated checksum. This command is only valid with packets utilizing the DDP long format (LAP Type field \$2).

Preparing a Packet

When you press the edit button for a particular packet in the main window, the edit window of Figure D-2 will appear and you will be shown the information of that packet. This window is divided into two main sections: the header and the data, with 18 editing fields. Only one editing field is active at a time. This is indicated by highlighting that field's rectangular box. There are several circular buttons, check boxes and command buttons (**OK**, **CANCEL** and **CLEAR**) used in preparing the packet. The standard Macintosh editing features apply to most of these controls. Some, however, need further clarification. These are:

- Pressing the **TAB** key causes Poke to verify the information in the current field before activating the next field. The same is true if you press the **RETURN** key (except within the packet's data field). If an error is detected while verifying a field, a beep will sound and Poke will return you back to the error's location. (Possible errors are described at the end of this section.)
- Clicking the mouse on a different editing field will verify the information in the currently active field. If there are no errors, Poke moves to the field clicked on.
- You may type data beyond that visible in the field. Leading blanks are automatically removed in the packet header fields.

Entering the Packet's Name

The packet's name is used only to visually distinguish the various packets from others in the main window. It may contain any sequence of printable characters, but it is suggested that you limit the number of characters to 16.

Entering Information in the Header Fields

Information in the packet header fields can be entered in any one of three ways:

<u>Decimal</u> :	Type in the digits (e.g. 128). This is the default entry type.
<u>Hexadecimal</u> :	All hexadecimal (hex) numbers are preceeded by a dollar sign (e.g., \$80 = 128).
<u>Binary</u> :	Binary numbers are preceeded by a percent sign (e.g., %1111 = \$0F = 15).

Leading zeros are ignored. When a field has been verified, the number entered is automatically converted to hex format.

Possible Error Conditions

- Value in field is out of range. (see AppleTalk Protocol documents for the permissible ranges of the various fields)
- Unknown character in field. Valid digits for decimal format are [0..9] (where this represents a range from zero to nine); valid digits for hex format is [0..9, a..f, A..F], and valid digits for binary numbers are [0,1].

Entering Packet Data Information

The following format must be followed when entering information into the packet data.

Data bytes can be entered into the packet in two ways: by typing in the ASCII character corresponding to the byte's value or by entering the byte's value in its hex form.

To enter the hex form, type a "\$" followed by the two digit hex number (e.g. \$84,\$01). Note that "\$1" is invalid, you must enter "\$01". Byte's whose value corresponds in the ASCII code to a graphic character can be entered by just typing in that character. Example: to enter a byte with the value "\$62", type "b"; for "\$42" type "B"; for "\$31" type "1". Other examples can be found in Figures D-2 and D-3.

Note: Since the dollar sign (\$) is a special character, you can only enter it in its hex form "\$24".

Poke will detect errors from the end of the data back to its beginning.

Editing Buttons

Various buttons in the edit window control the information that constitutes the packet. Each set of buttons is described below:

Packet Type: ☐ LAP ☐ DDP ☒ ATP

The Packet Type buttons are used to choose the header type as described in the protocols document. After clicking on a button, only the fields appropriate for that protocol type will be shown. The default is **ATP**. Only one button may be selected at a time.

☐ Req ☒ Rsp ☐ Rel

These three buttons are only used for an ATP packet. They are used to format an ATP request, response or release packet. The default is **Req**. As above, only one button may be chosen at a time.

☐ HO ☐ EOM ☐ STS

Each of these check boxes represents the corresponding bit in the ATP control field. If checked, the corresponding ATP control field bit will be set; otherwise the bit is cleared.

Packet Data Display
☐ Hex ☒ ASCII

The Packet Data Display buttons allow the user to select the type of display for the packet's data: hex strings or mixed ASCII and hex.

Note: This operation may take up to 10 seconds for large packets.

If an error occurs during the format conversion, an error message is displayed and the conversion will abort. You may enter data in either format at any time. The above buttons are used only when the display is updated or when you wish to convert data to the format immediately.

OK

CANCEL

CLEAR

The **OK** button should be pressed when you are through editing the packet. All fields are verified for correctness and the packet length is displayed before returning to the main window. You will also have the option, at this time, to calculate a checksum for the packet. If any errors are detected, you will be returned to the edit window.

The **CANCEL** button terminates the editing session without saving any changes to the packet. The packet is returned to the original form that it had prior to this editing attempt. Poke returns you to the send window.

The **CLEAR** button clears all editing fields and inserts the default information into them.

Sending Packets

To send packets, Poke must be in the send state (i.e., displaying the main window). Any one of the ten packets may be sent by clicking on its active **SEND** button. The number of times the packet will be sent and the delay between each of these transmissions is shown at the top right corner of the main window in the short form:

$$\text{Rpt Factor} = \text{nx} : \text{d ticks}$$

where: **n** = number of transmissions
 d = time interval between transmissions (in ticks)

If a **SEND** button is inactive, you must first edit the packet. The result of the **SEND** operation is displayed in the message area at the bottom of the main window.

Possible Error Conditions

- No error; packet was sent to destination node (or broadcast)
- -95; Packet was unable to be sent because either the destination node did not respond or the line was sensed "in use" 32 times.

Startup Notes

When Poke starts up, the MPP driver is opened and initialized. If the open call fails and you are returned a -35 error, you will be forced to exit the program. Most likely the cause of this error will be that the MPP driver is not installed in the System resource file. In addition, if the system heap is fragmented such that the MPP driver cannot get enough memory to load, the same error will be returned.

If the serial port configuration byte (SPConfig) is not set correctly, you will get a -98 error when Poke starts. See the AppleTalk Manager manual for additional information on location and contents of this byte.

Caveats

- Editing of the packet's data field will slow down appreciably as its size increases. Whenever possible, display it under the ASCII mode to minimize the number of screen characters.
- While in the ASCII display, all characters in the printable ASCII range (\$20-\$7E) and **RETURN** (\$0D) will be displayed in their ASCII form, even if they were entered as hex strings.

- The packet data field is limited to 55 lines. Even short packets (e.g., entering more than 55 carriage returns in the packet's data field in ASCII mode) can go out of the scrolling range.
- Numbers cannot be entered into the packet's data field in decimal or binary format.
- In no case can the size of the packet be greater than 603 bytes, including ALAP header.
- If an error occurs while verifying or converting the packet's data field, the information at the error location may change, as Poke tries to back out of the error gracefully.
- If you have chosen DDP or ATP packet types from the edit window, DDP long format will always be displayed, even if the ALAP type of \$01 (short format) was entered.
- If you enter more than 600 bytes of packet data, the checksum calculation may not work correctly until you have exited and reentered the editing window. (This will truncate all excess data from the end of the packet).



ULISSE s.r.l.
Servizio Distribuzione Apple
Via Pacini, 22
20131 Milano
Tel. (02) 2367783-2367406

su licenza esclusiva
Apple Computer S.p.A.

AppleTalk Filing Protocol

Version 1.1

IAT02

YOU SHOULD CAREFULLY READ THE FOLLOWING TERMS AND CONDITIONS BEFORE USING THIS SOFTWARE. ANY DOWNLOADING, REPRODUCTION, COPYING OR OTHER USE OF THE SOFTWARE WILL CONSTITUTE ACCEPTANCE OF THESE TERMS AND CONDITIONS.

SINGLE-COMPUTER END USER SOFTWARE LICENSE AGREEMENT

APPLE COMPUTER, INC. ("Apple") provides this software and licenses its use. You assume responsibility for the selection of the software to achieve your intended results, and for the installation and use of, and results obtained from, the software.

LICENSE

Pursuant to this license you may:

1. Use the software only on a single Apple computer. You must obtain a supplementary license from Apple before using the software in connection with systems and multiple central processing units, computer networks or emulations on mainframe or minicomputers.
2. Download the software only on media that is compatible with Apple manufactured computers.
3. Copy the software into any machine readable form for backup purposes in support of your use of the software on the single Apple computer.
4. Transfer the software and license to another party with a copy of this Agreement provided the other party reads and agrees to accept the terms and conditions of this Agreement. If you transfer the software, you must at the same time either transfer all copies, whether in printed or machine-readable form, to the same party or destroy any copies not transferred. Apple grants a license to such other party under this Agreement and the other party will accept such license by its initial use of the software. If you transfer possession of any copy of the software, in whole or in part, to another party, your license is automatically terminated.

This software is protected by United States copyright law. You must reproduce the Apple copyright notice on any copy of the software.

THIS SOFTWARE MAY BE ELECTRONICALLY DISTRIBUTED ONLY BY AUTHORIZED ELECTRONIC DISTRIBUTORS. IT MAY BE DOWNLOADED ONLY FOR PERSONAL OR NON-COMMERCIAL USES ON APPLE COMPUTERS AND MAY NOT BE REDISTRIBUTED OR USED FOR COMMERCIAL PURPOSES WITHOUT AN EXPRESS SOFTWARE DISTRIBUTION LICENSE FROM APPLE. These licenses are available from Apple's Software Licensing Department.

YOU MAY NOT MODIFY, REVERSE COMPILE, DISASSEMBLE, NETWORK, RENT, LEASE, LOAN OR DISTRIBUTE THE SOFTWARE, OR ANY COPY, IN WHOLE OR IN PART. YOU UNDERSTAND THAT UNAUTHORIZED REPRODUCTION OF COPIES OF THE SOFTWARE OR UNAUTHORIZED TRANSFER OF ANY COPY OF THE SOFTWARE MAY SUBJECT YOU TO A LAWSUIT FOR DAMAGES, INJUNCTIVE RELIEF, AND ATTORNEY'S FEES.

Apple reserves all rights not expressly granted to you.

Export law assurances

You agree and certify that neither the software and documentation nor any direct product thereof (1) is intended to be used for nuclear proliferation or any other purpose prohibited by the United States Export Administration Act of 1979, as amended (the "Act") and the regulations promulgated thereunder, and (2) is being or will be downloaded, shipped, transferred or reexported, directly or indirectly, into any country prohibited by the Act and the regulations promulgated thereunder.

Government End Users

If you are acquiring the software on behalf of any unit or agency of the United States government, you agree that: (a) the software is "Commercial Computer Software" as that term is

defined in Paragraph 27.401 of the DoD Supplement to the Federal Acquisition Regulations (the "Supplement") or is within the equivalent classification of any other federal agencies' regulations; (b) the software was developed at private expense, and no part of it was developed with government funds; (c) the government's use of the software is subject to "Restricted Rights" as that term is defined in clause 52.227-7013 (b)(3)(ii) of the Supplement or in the equivalent clause of any other federal agencies' regulations; (d) the software is a "trade secret" of Apple for all purposes of the Freedom of Information Act; and (e) each copy of the software will contain the following Restricted Rights Legend:

"Restricted Rights Legend"

Use, duplication or disclosure is subject to restrictions as set forth in subdivision (b)(3)(ii) of the Rights in Technical Data and Computer Software clause at FAR 52.227-7013. Manufacturer: Apple Computer, Inc. 20525 Mariani Avenue, Cupertino, California 95014.

You agree to indemnify Apple for any liability, loss, costs and expense (including court costs and reasonable attorneys' fees) arising out of any breach of the provisions of this Agreement relating to use by the government.

Term

The license is effective until terminated. You may terminate it at any time by destroying the software together with all copies. The license will also terminate upon conditions set forth elsewhere in this Agreement or if you fail to comply with any of the terms or conditions of this Agreement. You agree upon such termination to destroy all copies of the software.

Disclaimer of Warranty

The software is provided "as is" without warranty of any kind, either express or implied, with respect to its merchantability or its fitness for any particular purpose. The entire risk as to the quality and performance of the software is with you. Should the software prove defective, you (and not Apple or an Apple authorized representative) assume the entire cost of all necessary servicing, repair or correction.

Apple does not warrant that the functions contained in the software will meet your requirements or that the operation of the software will be uninterrupted or error free or that defects in the software will be corrected.

Some states do not allow the exclusion of implied warranties, so the above exclusion may not apply to you. This warranty gives you specific legal rights and you may also have other rights which vary from state to state.

Limitation of Remedies

In no event will Apple be liable to you for any lost profits, lost savings or other incidental, special or consequential damages arising out of the use of or inability to use any software even if Apple or an authorized Apple representative has been advised of the possibility of such damages, or for any claim by any other party.

Some states do not allow the limitation or exclusion of liability for incidental or consequential damages so the above limitation or exclusion may not apply to you.

Apple's liability to you for actual damages for any cause whatsoever, and regardless of the form of the action, will be limited to the greater of \$500 or the money paid for the software that caused the damages or that is the subject matter of, or is directly related to, the cause of action.

General

This Agreement, if any attempt to network, rent, lease, or sublicense the software, or, except as expressly provided in this Agreement, to transfer any of the rights, duties or obligations under this Agreement, becomes void.

The Agreement will be construed under the laws of the state of California, except for that body of laws dealing with conflict of laws. If any provision of this Agreement shall be held by a court of competent jurisdiction to be contrary to law, that provision will be enforced to the maximum extent permissible, and the remaining provisions of this Agreement shall remain in full force and effect.

AppleTalk Filing Protocol (AFP) Engineering Technical Notes

Apple Computer Inc.

Protocol Version 1.1
February 17, 1987

This protocol is proprietary and based on AFP Version 1.0 which was developed in joint work by Apple Computer Inc. and Centram Systems West

Table of Contents

Contents

1	Chapter 1: AppleTalk Filing Protocol Design Description
1	The Basic File Access Model
3	Goals of AFP
3	AFP in the AppleTalk Architecture
5	Notation
6	Chapter 2: AppleTalk Filing Protocol System Description
6	AFP File System Structure
6	File Servers
8	Volumes
8	The Volume Catalog: Directories, Files, and Forks
10	CNode Names: Long and Short Names
11	Directory IDs
11	Volume Signatures
12	Directory and File Parameters
14	Date-Time Values
14	Pathnames
15	Access Paths and Open File Forks
15	Complete Specification of a Catalog Node (CNode)
17	Chapter 3: The AFP Login Procedure
17	AFP Versions
17	User Authentication Methods

18	Discovering a File Server:
18	Obtaining File Server Information:
18	The Login Step:
19	Chapter 4: Access Control Mechanisms in AFP
19	User Authentication at Server Login
19	Volume Passwords
19	Directory-Level Access Controls
23	Chapter 5: A Discussion of AFP Calls
23	Server-Level Calls
24	Volume-Level Calls
24	Directory-Level Calls
25	File-Level Calls
25	Combined Directory-File-Level Calls
26	Fork-Level Calls
27	Chapter 6: A Design for the Finder's Desktop Management in a Network Environment
31	Chapter 7: Specification of AFP Calls
32	FPAddAPPL
33	FPAddComment
34	FPAddIcon
36	FBByteRangeLock
38	FPCloseDir
39	FPCloseDT
40	FPCloseFork

41	FPCloseVol
42	FPCopyFile
44	FPCreateDir
45	FPCreateFile
47	FPDelete
48	FPEnumerate
51	FPFlush
52	FPFlushFork
53	FPGetAPPL
54	FPGet Comment
55	FPGetFileDirParms
58	FPGetForkParms
59	FPGetIcon
60	FPGetIconInfo
61	FPGetSrvrInfo
63	FPGetSrvrParms
64	FPGetVolParms
65	FPLogin
67	FPLoginCont
68	FPLogout
69	FPMaPID
70	FPMaPName
71	FPMove
73	FPOpenDir
74	FPOpenDT
75	FPOpenFork
77	FPOpenVol
79	FPRead
81	FPRemoveAPPL
82	FPRemoveComment
83	FPRename
85	FPSetDirParms
88	FPSetFileParms
90	FPSetFileDirParms
93	FPSetForkParms
94	FPSetVolParms
95	FPWrite

97 **Chapter 8: AFP's Use of ASP**

99 **Appendix A: User Authentication Methods**

99	No User Authentication
99	User Authentication With Cleartext Password Transmission
99	User Authentication Based on Random Number Exchange

101	Appendix B:	Long/Short Name Management Algorithms
103	Appendix C:	File Sharing Modes
105	Appendix D:	Values of Command and Error Codes
107	Error Codes	
108	Appendix E:	List of Abbreviations

Chapter 1

AppleTalk Filing Protocol Design Description

This document specifies version 1.1 of the *AppleTalk Filing Protocol* (AFP). We start by discussing, in general terms, the scope of this protocol and some strategic decisions that have had a fundamental impact on its design. The relationship of this protocol to other parts of the AppleTalk protocol architecture are then spelled out. The detailed specification of the protocol is presented in three main parts: the AFP system model, AFP calls, and AFP packet formats.

The Basic File Access Model

The overall objective is to allow workstations on AppleTalk to access files on file servers connected to the network. This access occurs at the level of calls to the workstation's native file system. Figure AFP1 illustrates the basic file access model which can be used to make these general notions more precise.

As illustrated, a program running in a workstation issues *native file system commands*. If these commands refer to files on the workstation's local storage, then they are processed by the workstation's local/native file system. If however the commands refer to files on a file server, they are routed inside the workstation to a *Translator*. It is the Translator's responsibility to deliver the desired file system service to the requesting program. This is done by converting the native file system command into one or more AFP requests sent over the network to the appropriate file server. We can isolate two important file service interfaces in Figure AFP1.

The first, labeled *Native Filing Interface* (NFI), is the one through which programs on the workstation issue native file system commands and obtain the relevant file system services. A program "looking through" this interface sees the structures and capabilities of the workstation's native file system.

The second interface, labelled *AppleTalk Filing Interface* (AFI), is the one through which the Translator issues AFP requests to the file server and obtains AFP services. Looking through this interface, the Translator sees what we will call the AFP file system. It is the Translator's responsibility to map requests made through NFI into the requests that must be made through the AFI to the AFP file system in order to satisfy the original NFI request.

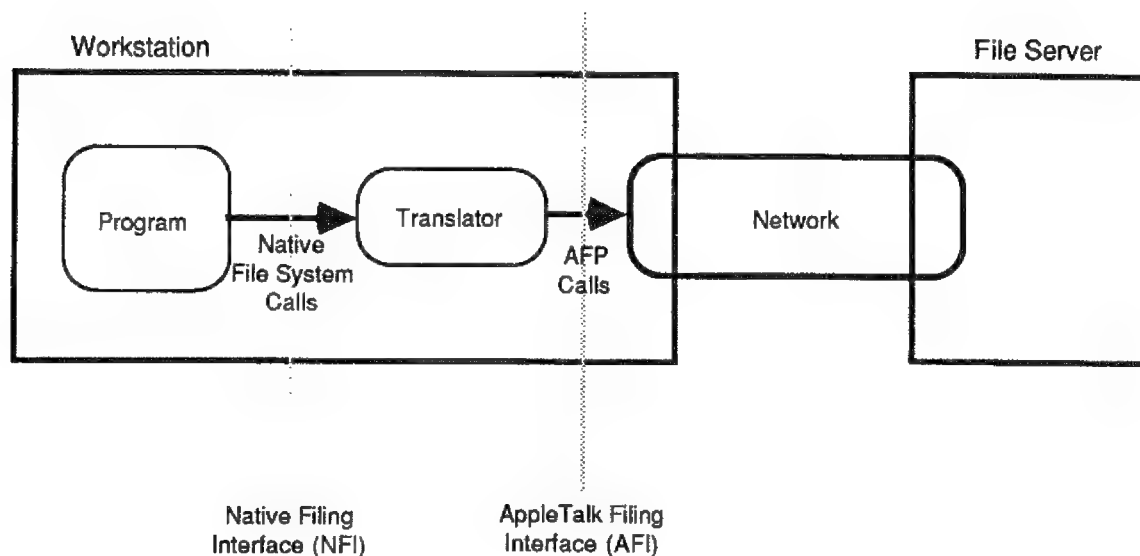


Figure AFP1. The Basic File Access Model

Note that in some cases it may be necessary to provide another path from the Program directly to the AppleTalk Filing Interface to allow the Program to make AFP calls which have no equivalent in the Native Filing Interface.

The filing protocol specification in effect is the same as a complete specification of the AFI. This consists of three parts:

- the AFP file system structure;
- the AFP calls;
- the algorithms associated with these calls.

By the AFP file system structure is meant a complete description of entities such as servers, volumes, directories, files, forks, etc., that are "visible" through the AFI, their interrelationship and associated parameters. The client of the AFI (the Translator) sends calls/commands/requests through the AFI to manipulate this AFP file system structure, and the details of what each call does to this structure constitute the algorithms associated with AFP.

It is interesting to consider the relationship between the native file system and the AFP file system. Clearly, the latter must be functionally as powerful as the former. In fact, when the AFP is designed to allow various different types of workstations to use the model of Figure AFP1 at the same time, then the services provided by the AFP file system must, in a sense, be the functional union of the services provided by the various native file systems of these workstations.

This document does not examine in detail the mechanism of translation in the Translator; the focus here is on the AFI and the AFP itself.

Goals of AFP

We have paid special attention in the design of AFP to allow its extension in a very general fashion. This is essential if, in the future, additional types of workstations are to be supported by the protocol.

AFP version 1.1 is designed to work with Macintoshes (using the hierarchical file system) and MS-DOS workstations. Thus, the protocol is sufficient to successfully support translators on these two workstations under the indicated operating environment. We expect *future* extensions of the protocol to support the Apple-II (under ProDOS) and Unix workstations (but these are not part of the goals of version 1.1).

Access control mechanisms are an important part of filing protocols. AFP supports user authentication in a flexible fashion allowing the use of two standard password-based schemes but permitting future extension to other more sophisticated methods based, for instance, on public key encryption. AFP does not force the use of a fixed user-authentication method. AFP includes a directory-level access control mechanism based on user authentication at server login time. This scheme is similar to that used by Unix and therefore should be especially easy to implement on Unix-based server machines.

Although in the model of Figure AFP1 we have distinguished between workstations and file servers, AFP does not rule out the possibility of a particular network node being both a workstation and a file server. At the same time, AFP does not attempt to solve various concurrency problems and potential deadlock situations that can arise in such combined workstation/server nodes. This must be done by careful design of the software implemented on such "combined" nodes.

It should also be mentioned that AFP does not include any services or calls needed for administration of file servers. These will depend on the nature of a particular file server and are hence outside the scope of AFP.

AFP in the AppleTalk Architecture

As illustrated in Figure AFP2, AFP is a client of the AppleTalk Session Protocol (ASP) described in a separate document.

As should be clear from the ASP specification, a file server must call the AppleTalk Transaction Protocol (ATP) to open up a Session Listening Socket (SLS) and then use the Name Binding Protocol (NBP) to register the file service's name on this socket. Workstations wishing to use the file server must use NBP and the file service's name to discover the SLS's network address. With this information the workstation opens an ASP session with the server.

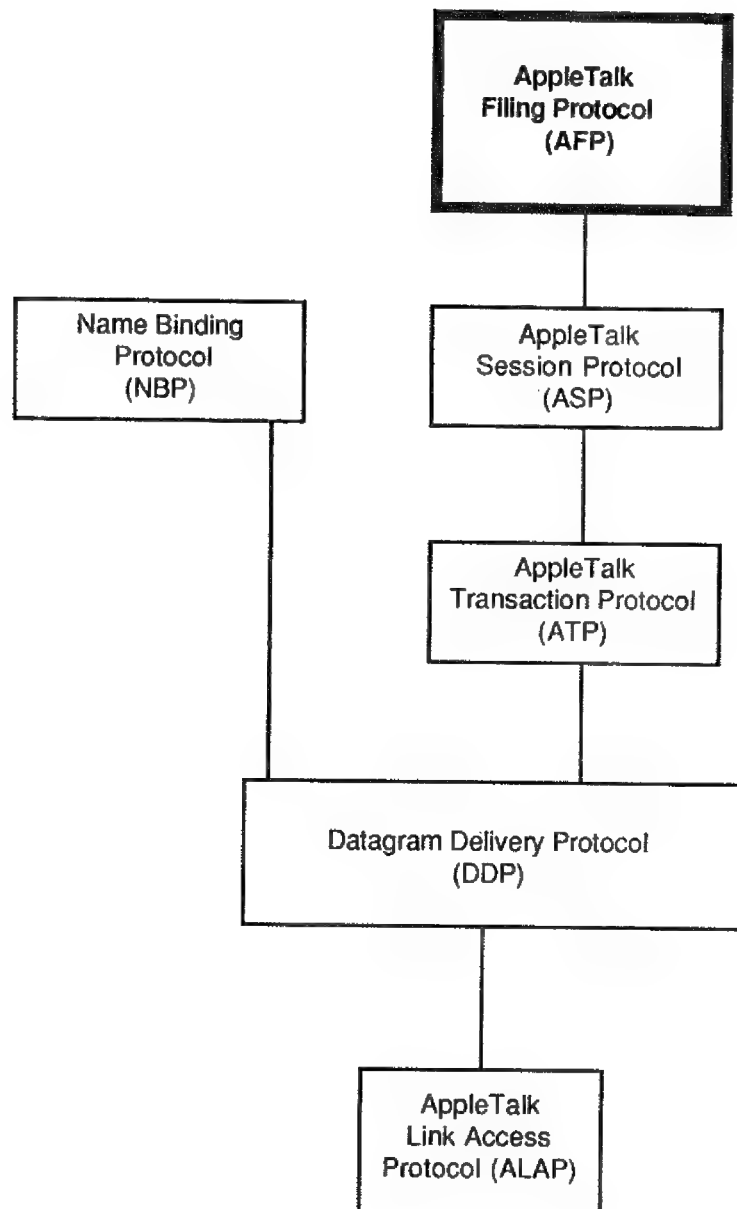


Figure AFP2. AFP and the AppleTalk Protocol Architecture

Once this session has been opened the AFP protocol entity in the workstation must log itself in on the file server as a bonafide user. The login step provides an opportunity to:

- authenticate the file server's user (i.e., the workstation);
- negotiate the version of the AFP protocol to be used during that session.

After the workstation has successfully logged in on the file server, the various AFP commands can be conveyed on the session thus established. When the workstation has finished using the services of the file server, it can logout. At this point all resources related to this filing session are freed up in the file server, and the underlying ASP session is closed.

Notation

Throughout this document hexadecimal (base 16) values are written with a leading \$ sign (for example, \$3A), while decimal integers are written with no leading special character (for example, 18).

Acronyms are widely used in this document (a list is provided in Appendix E).

The following abbreviations are used to describe the input and output parameters of AFP calls:

BIT	a single binary digit
BUF	a buffer; exact method of specifying location and length are dependent on the AFP implementation
BYTE	an 8-bit quantity
EntityAddr	a network-visible entity's internet address; exact size and format are dependent on the underlying network
INT	a 2-byte (16-bit) integer quantity
LONG	a 4-byte (32-bit) quantity
STR	a string consisting of a one-byte string length value (not including the length byte) followed by the string's characters (one character per byte). Strings cannot have more than 255 characters.
ResType	a 4-byte signature used in Finder info fields

Note that all string comparisons in AFP are case insensitive unless otherwise noted.

All numerical quantities are signed numbers unless otherwise noted.

Chapter 2

AppleTalk Filing Protocol System Description

This chapter contains a fairly detailed discussion of all the major concepts involved in AFP: the AFP file system structure, the login process, access control mechanisms provided by AFP, and a global discussion of the AFP calls.

Chapter 7 contains the detailed specification of the AFP calls. This forms the bulk of this document.

Although AFP has been designed to be used with various underlying transport mechanisms, on AppleTalk it is implemented as a client of the AppleTalk Session Protocol (ASP). In Chapter 8 we include a discussion of how AFP uses the services provided by ASP.

AFP File System Structure

Once the workstation client of AFP has logged into the server, it can issue any of a set of AFP calls described below. Through these calls the client can obtain descriptive information (collectively referred to below as parameters) about any of a set of entities that comprise the *AFP file system structure*, modify this descriptive information, create or delete entities, read and write to such entities, etc. Before discussing the calls for manipulating the AFP file system structure let us discuss this structure, its entities as seen through the AFI, and the relationship among these entities. [Note that this discussion does not describe how this structure might be implemented on a server; it just describes the final result of such an implementation.]

The entities comprising the AFP file system structure are: [file] servers, volumes, directories, files, and forks. We now examine each of these and the relationship among them.

File Servers

The most global entity encompassed by the AFP file system structure and visible through the AFI is a *file server* (*server* for short). File servers possess names and can be discovered by a workstation as discussed below in connection with the login process. Servers manage one or more volumes which can be accessed through the AFP. (See Figure AFP3.)

There are several parameters associated with a server that are visible through the AFI:

- *server name* [string: max 32 characters]
- *server machine-type identification* [string: max 16 characters]




Figure AFP3. File Servers and Volumes

- *number of volumes* on the server [2-byte integer]
- *AFP version strings* [strings: maximum 16 characters per string]
- *UAM strings* [strings: maximum 16 characters per string]

In addition, a server can maintain the following optional parameter, which can be used to customize the appearance of a server's volumes on a Macintosh workstation's desktop:

- *server icon* [256 bytes].

The server icon consists of a 32-by-32 bit (128 bytes) icon bitmap followed by a 32-by-32 bit (128 bytes) icon mask. The mask usually consists of the icon's outline filled with black (a bit that is set). This format is exactly that required for the specification of icons for a Macintosh (see the *Structure of a Macintosh Application* chapter of *Inside Macintosh*).

The *server machine-type identification* is provided through a string of at most 16 characters. This string is purely informative, providing a textual description of the type of hardware and/or software system on which the file server has been implemented; it has no significance as far as the AFP is concerned.

The *AFP version strings* and the *UAM strings* will be discussed below in connection with the login process.

Volumes

As noted a server manages one or more *volumes* that are visible to workstations through the AFI. For each volume the server must maintain and make visible through the AFI the following parameters:

- *volume name* [string: max 27 characters]
- *volume signature* [2 bytes]
- *volume identifier* (VolID) [2 bytes]
- *volume creation date-time* [4 bytes]
- *volume modification date-time* [4 bytes]
- *volume backup date-time* [4 bytes]
- *volume size in bytes* [4-byte unsigned integer]
- *number of free bytes on volume* [4-byte unsigned integer].

In addition a server can maintain the following optional volume parameter which can be used to provide a simple kind of security at the volume level:

- *volume password* [8 bytes].

The *volume name* is a string of up to 27 bytes which identifies a server volume. The permissible characters in a volume name are all 8-bit characters not including null (\$00) or colon (\$3A). A given server's volumes must all have different names. The volume name is used only to identify the volume to the workstation's user. It is not used directly as part of the specification of objects on the volume (see the discussion of pathnames below). Instead, the workstation client makes an AFP call to obtain a particular volume's VolID, which is then used as the volume's identifier in all subsequent AFP calls.

The *volume signature* is a 2-byte quantity used to provide a volume-type identifier. In version 1.1 of AFP only three values of the volume signature are permitted (see the discussion below of volume signatures).

Within each filing session the server assigns a *volume identifier* (VolID) to each of its volumes. This value is unique (within a particular filing session) among the volumes of a given server and can be used in the AFP calls sent over a filing session to uniquely identify the volume to which the calls applies.

The *volume creation*, *modification*, and *backup date-time* quantities are maintained in terms of GMT values (see the discussion below). A volume's creation date-time is set by the server when the volume is created. Likewise, the modification date-time is changed by the server every time the volume is modified (see the *FPFlush* call). These two date-time values are managed solely by the server and they cannot be modified by the AFP client. However, the backup date-time is meant to be appropriately set by a backup program each time the volume's contents have been successfully backed up. When a volume is first created, its backup date-time is set to \$80000000 (the earliest representable date-time value).

Directories and files are stored in volumes and constitute the next level of entities visible through the AFI.

The Volume Catalog: Directories, Files and Forks

A volume consists of directories and files arranged in a hierarchical structure known as the *volume catalog* (see Figure AFP4).

The *volume catalog* (also referred to as the *catalog*) is a description of the contents of the volume organized as a tree. The nodes of this tree, known as *catalog nodes* (*CNode* for short) are either files or directories. The *internal* (non-terminal) nodes of the tree are always directories, while the *leaf* (terminal) nodes are files or empty directories. At the base of each catalog is a special CNode (a directory) called the *root*.

A catalog does not span multiple volumes; the AFP client will see a separate catalog for each server volume visible through the AFI.

Directories should be looked upon as "logical containers" which contain other directories and/or files. Thus directories are equivalent to the concept of folders in the Macintosh user interface.

As in the Macintosh file system, a file consists of two *forks*: one, the *data fork*, is an unstructured finite sequence of bytes; the other, the *resource fork*, is typically used to hold Macintosh OS resources and a data structure for mapping them within the fork. As far as AFP is concerned both forks are simply finite-length byte sequences; specification of the resource structure of the resource fork is outside the scope of this protocol. Note that the bytes in a file fork are numbered starting with zero (the first byte in the fork is numbered 0).


A given file can have either or both forks empty. In fact, files created by an MS-DOS machine will most likely have an empty resource fork since this construct is outside the conceptual structure of an MS-DOS file and hence unintelligible to such a system. Likewise an MS-DOS machine accessing a server file created by a Macintosh will, in all likelihood, not access the resource fork of the file, and will in fact be unaware of its existence and significance. However, if applications are written for MS-DOS machines which understand and manipulate resource forks, they may do so; AFP does nothing to prevent this.

Non-Macintosh-based AFP clients that need only one file fork must use the data fork. It is important that the internal structure of the resource fork in terms of resources be maintained correctly. This is essential, since that is what a Macintosh workstation expects to see in the resource fork of any file. For this reason, workstations that do not know how to manage the internal structure of the resource fork should never alter its contents.

The tree structure of a catalog leads naturally to the idea of a CNode's *parent*, i.e., the node containing the given CNode. In terms of the pictorial representation of Figure AFP4, a given CNode's parent is the CNode immediately above it in the catalog tree. A CNode's parent will often also be referred to as its *parent directory*.

The set of CNodes contained in a given directory will be referred to as its *offspring*. This possibly empty set (for an empty directory) is exactly the set of CNodes for which the given directory is the parent. The *valence* of a directory indicates the number of offspring it contains.

Note that the root of the catalog is special in that it has no parent. For a given CNode *X* there is one path from the root to that CNode. Any CNode on that path is known as an *ancestor* of *X*. *X* is said to be a *descendant* of any of its ancestors.



Every CNode (including the root) has two names: a *long name* and a *short name*.

Short names are strings of up to 12 characters that follow the well-known 8.3 format (<8 characters>.<3 character extension>) of file/directory names used by the MS-DOS operating system, and consist of a sequence of up to 8 characters followed by an optional

extension field consisting of a period (\$2E) and up to 3 characters. All characters must be alpha-numeric or from the set ! # \$ % &) (, - @ _ { } ~ (period is not included).

It should be clear why AFP attaches two names to each CNode, one each for the two native file systems supported by version 1.1 of AFP. Macintosh workstations expect to refer to files and directories by long names, while MS-DOS machines will use the 8.3 format short names. By using appropriate algorithms for creating the "other" name (long or short) when creating or renaming files and directories, a consistent sharing of these objects is possible from these dissimilar workstations.

The rules of uniqueness of long and short names are quite simple. No two offspring of a given directory can have the same short name or the same long name. A short name may match a long name if they both belong to the same object. Thus either name (long or short) uniquely identifies CNodes within the context of a given parent. To ensure that this uniqueness is maintained at all times, a file server must implement carefully-designed file/directory creation and renaming algorithms discussed in the specification section of this document and in Appendix B.

The root of the catalog is in fact a "container" representation of the volume. The root's long name is exactly the same as the volume name, so is therefore restricted to a maximum size of 27 bytes. The root of a volume can neither be deleted nor renamed through AFP. Hence a volume cannot be renamed through AFP.

Directory IDs

Each directory in the catalog has in addition to its names a numerical (4-byte) identifier known as its *directory ID* (*DirID* for short). The DirID of the root is always equal to 2.

Two directories on a given volume cannot have the same DirID. Thus the DirID uniquely identifies a directory on a given volume.

The DirID of a given CNode's parent is said to be the CNode's *ParentID*. This is a special case of the more general term *AncestorID*, which is the DirID of a CNode's ancestor. Although a given CNode has a unique ParentID (there is only one parent), it may have several AncestorIDs (one for each ancestor). The root's Parent DirID is always equal to 1. Zero is not a valid DirID.

Volume Signature

AFP allows three types of volumes. To determine what type a particular volume is, we associate a volume signature with each volume. This is a 2-byte quantity that can, for AFP version 1.1, take one of three permissible values:

- 01 Flat volume
- 02 Fixed DirID volume
- 03 Variable DirID volume.

Any value other than these should be taken as referring to a volume unintelligible to AFP version 1.1 and is hence an error condition.

A *flat volume* is one whose catalog tree consists of only one directory, the root, containing files. As an example consider the flat Macintosh file system's flat volumes. Attempting to

create a directory on a flat volume results in an error. The specification of files on such a volume consists of the VolID, a DirID having a value of 2, and the file name.

A *fixed DirID volume* is one that has a hierarchical catalog in which the directory IDs of the various directories in the catalog are fixed and determined at the time of creation of a directory. The DirID of a directory never changes during the lifetime of the directory. Furthermore, this DirID will not be used for any other directory during the lifetime of the volume (i.e., it is not used again even if the corresponding directory is deleted from the volume).

A *variable DirID volume* differs from a fixed DirID volume in that, although it maintains the uniqueness of DirIDs, it does not associate a fixed DirID with a directory over the entire lifetime of the directory. In fact, for such volumes the file server creates a unique DirID for a directory when the workstation client issues an FPOpenDir call (discussed in Chapter 7) and maintains this DirID value of a directory either until an FPCloseDir call is issued by the client or until the AFP session is terminated. Using a DirID obtained through an FPOpenDir call outside of this range of time will lead to unpredictable results.

Clearly, MS-DOS machines can access any of the three types of server volumes since the concept of DirIDs is foreign to their file system. However, it should be noted that the Macintosh using either MFS or HFS cannot directly use variable DirID server volumes (volume signature = 03). Macintosh HFS volumes are fixed DirID volumes and Macintosh applications such as the Finder save the DirIDs of various directories and expect these to remain invariant; hierarchical catalog volumes can be handled by this file system (HFS) only if they are fixed DirID volumes (volume signature = 02).

It should be noted that particular directories of a variable DirID volume can be mounted as flat MFS "volumes" and used by a Macintosh workstation. In this case, the corresponding "virtual" volume will be seen as flat in structure and any directories contained in that "volume"/directory will not be accessible through AFP. This is not a recommended way of implementing file servers for the Macintosh and its use is discouraged. In fact variable DirID volumes are included in our definition of AFP only to allow non-Macintosh machines (whose file systems are unable to implement, in an easy way, the fixed DirID feature) to function as simple servers offering files within particular directories for access over AppleTalk.

Directory and File Parameters

A server must maintain the following parameters for each directory:

- *long name* [string: max 31 characters]
- *short name* [string: max 12 characters]
- *DirID* [4 bytes]
- *ParentID* [4 bytes]
- *attributes* [2 bytes]
- *Finder Information* [32 bytes]
- *number of files and directories contained in the directory* [2-byte integer]
- *creation date-time* [4 bytes]
- *modification date-time* [4 bytes]
- *backup date-time* [4 bytes]
- *owner ID* [4 bytes]

- *group ID* [4 bytes]
- *owner's Access Rights* [1 byte]
- *group's Access Rights* [1 byte]
- *world's Access Rights* [1 byte].

The *Finder Information* is exactly the Finder information for directories that must be saved for HFS directories. This field is not examined by AFP. It is written (set) by the client of AFP. The last five fields in this list are related to access controls in AFP and will be discussed later in this document.

The *attributes* parameter is a bitmap indicating various attributes of the directory. For AFP version 1.1, one directory attribute is defined: *Invisible* (this directory should not be made visible to the workstation's user).

A server must maintain the following parameters for each file:

- *long name* [string: max 31 characters]
- *short name* [string: max 12 characters]
- *ParentID* [4 bytes]
- *file number* [4 bytes]
- *attributes* [2 bytes]
- *Finder Information* [32 bytes]
- *data fork length* [4-byte unsigned integer]
- *resource fork length* [4-byte unsigned integer]
- *creation date-time* [4-byte signed integer]
- *modification date-time* [4-byte signed integer]
- *backup date-time* [4-byte signed integer].

The *file number* is a 32 bit number associated with a given file, unique among all files on the volume. This is purely informative, since AFP does not allow the specification of a file by its file number.

The *attributes* parameter is a bitmap indicating various attributes of the file. For AFP version 1.1, five file attributes are defined. The rest of the eleven bits must be equal to zero. The five attributes are: *Read-Only* (cannot write to the file's forks), *DAIreadyOpen* (the file's data fork is currently open by some user), *RAIreadyOpen* (the file's resource fork is currently open by some user), *Multi-User* (the file is an application that has been written for simultaneous use by more than one user) and *Invisible* (this file should not be made visible to the workstation's user).

The *resource* and *data fork lengths* are equal to the number of bytes in the corresponding fork.

The *creation date-time* of a directory or a file is set to the value corresponding to the server's system clock when the file or directory is created. An AFP client with the appropriate access rights can set this date-time to any desired value. The *backup date-time* values are for the use of backup programs. When a file or directory is created the server sets the backup date-time to \$80000000.

The *modification date-time* of a file is changed by the server every time either of the file's forks is closed or flushed, if this fork has been written to in that filing session (see the *FPClose* call). Furthermore, an AFP client with the appropriate access rights can set this date-time to any desired value.

The *modification date-time* of a directory is changed by the server every time the directory's contents are modified: this includes renaming the directory, creating or deleting a CNode in the directory, moving the directory, or changing its access privileges. Furthermore, an AFP client with the appropriate access rights can set this date-time to any desired value.

Date-Time Values

All date-time quantities used by AFP are Greenwich Mean Time (GMT) values. They are 32-bit signed integers corresponding to the number of seconds measured from 12:00 am on January 1, 2000 (the start of the next century corresponds to date-time=0).

The use of GMT makes these AFP date-time quantities independent of the geographical location of the servers and workstations using AFP. The Translator in a given workstation must, if it wishes to present local time-date values, carry out the appropriate conversion based on the geographical time zone and date (to allow for the changes related to Daylight Savings Time, etc.) applicable to the workstation.

Simple implementations of AFP workstations may wish to just obtain the server's time at login (using the `FPGetSrvrParms` call), calculate the offset between it and the workstation's clock, and then add or subtract this offset to all date-times sent to or received from the server.

Pathnames

As noted above, CNodes (files and directories) have long and short names, either of which can be used to uniquely identify a CNode within the context of its parent directory.

CNode names can be concatenated with intervening null-byte separators to form *pathnames*. Each element of a pathname must be the name of a directory, except for the last one which can be the name of a directory or a file.

The elements of a pathname can be long or short names; yet a given pathname cannot contain a mix of long and short CNode names. With a given pathname it is necessary to associate a *pathname type* which indicates one of two situations: the elements of the pathname are all short names (*pathname type* = 1), or all long names (*pathname type* = 2).

As the term signifies, a pathname can be used to traverse the catalog tree. The starting point of this path must be a directory which is separately defined (by its DirID). The pathname must be parsed from left to right to obtain each element which is used as the next node on the path. A valid pathname must proceed step by step from parent to offspring. The first element of the pathname is interpreted as an offspring of the starting point directory. A single null-byte separator preceding this first element is ignored.

Provision is made in the pathname syntax for ascending from a particular CNode on the path up to its parent. This is indicated by two consecutive null-bytes in the separator. Three consecutive null-bytes in a separator mean move up two levels in the catalog tree, and so on.

The syntax of an AFP pathname can be summed up in the following BNF-like form:

`<Sep> ::= <null-byte>+`

`<Pathname> ::= empty-string |
 <Sep>* <CNode name> (<Sep><Pathname>)*`

where we use the notation (a*) to mean a sequence of zero or more (a)'s, while (a+) stands for a sequence of one or more (a)'s. From this BNF-like specification of pathnames it should be clear that it is a concatenation of CNode names delineated by separators consisting of one or more null bytes. Pathnames may also start or end with a string of null bytes.

AFP does *not* allow the inclusion of the root directory in pathnames, i.e., in the commonly used terminology of pathnames, AFP does not allow the use of full pathnames. The equivalent of a full pathname is achieved in AFP by specifying the starting point of the pathname as the root (DirID = 2) and presenting the pathname as descending down from the root (but not including the root).

Pathnames sent in AFP packets are formatted as Pascal strings starting with a length byte followed by up to 255 characters of the pathname (each pathname element does not include a length byte). In addition the pathname type must be provided. A single null byte is used to indicate that no pathname is supplied; note that even in this case, a valid pathname type must be provided.

Access Paths and Open File Forks

To read or write the contents of the data or resource fork of a particular file, the AFP client in the workstation must issue a special call to open the particular fork of the specified file. This leads to the creation of an access path to that file fork, and subsequent read and write calls will be processed by reference to that access path. To allow for this the server must generate an access path identifier known as the *Open Fork Refnum*, unique among all open forks on a given session. Also associated with a particular access path is an *Access Mode Descriptor*, which indicates whether this access path (open file fork) was opened for (and allows) reading and/or writing through that path.

For each access path (open file fork) the server must maintain the following parameters:

- *Open Fork Refnum* [2-byte integer] (zero is invalid)
- *Access Mode Descriptor* [2-byte bitmap]
- *Which-fork identifier* [1 significant bit in a byte].

In addition, the server must provide a way of accessing the file parameters of the file to which this open fork belongs.

The *Access Mode Descriptor* is maintained by the server and is inaccessible to workstation clients of AFP. It is used by the server to indicate the access mode with which that access path is open.

Complete Specification of a Catalog Node (CNode)

In AFP a particular CNode (file or directory) can be unambiguously specified to the server by providing the following: VolID, AncestorID, pathname. This specification subsumes special cases such as the specification of a particular directory just by VolID, AncestorID = DirID of the directory, and a null pathname (length byte of the pathname string equal to 0).

Chapter 3

The AFP Login Procedure

This chapter discusses the process of discovering a file server and logging in to it.

This process consists of three steps: (i) discovering the file server, (ii) obtaining file server information, (iii) the actual log on step. Before discussing these three steps we need to establish two relevant concepts: AFP versions and user authentication methods.

AFP Versions

As noted in the introduction, the AFP has been designed to be extensible. One of the techniques used towards this end is to associate one or more version descriptors, known as the AFPVersion with the protocol. These are strings of up to 16 characters which uniquely identify a particular version of the protocol. For instance for the version defined in the present document the AFPVersion string is "AFPVersion 1.1".

When logging in to a server the workstation must first find out which versions of AFP the server can handle. Then the workstation client of AFP selects the version it wants used on the session and indicates this in the session opening request.

User Authentication Methods

As part of the login process, AFP optionally provides for user authentication by the server. This involves conveying to the server some information identifying and validating the user. This can be done in many different ways depending on the level of security desired.

The basic problem is to identify the user to the server (this can be done with a user name) and then verify that this is a valid user. This user verification/authentication step can be based for instance on establishing in some way that this user has a secret piece of information (we will call this a password).

For instance, the user can send the user name and the password in clear text to the server which compares them against its list of valid information. The problem with this method is that anyone listening on the network with a "peek" utility can pirate this user/password information and use it for accessing the server.

A more secure method would be to send the server just the user's name. The server then consults its data base of user authentication information and extracts from it the user password. Then it uses a scheme based on a random number modification/encoding by the password to verify if the user in fact knows this password. The password is never sent over the network; in fact the encoding is done in such a way that the password cannot be extracted from the information transmitted over the wire.

In the future we expect many other methods to be developed for authenticating users over a network. To make AFP extensible in this respect and to allow different servers to choose

to implement any or all of these user authentication methods AFP expects the workstation client to obtain a list of all the user authentication methods the server handles, and then select one of these methods for the login step. The identification of user authentication methods is done by a string of up to 16 characters known as UAM. For AFP version 1.1 we describe three standard methods identified by UAM = "No User Authent", "Cleartxt passwd", and "Randnum exchange". Details of these three methods are discussed in Appendix A.

Discovering a File Server

As noted above, a file server implementing AFP must, upon being started up, open a session listening socket (SLS) and register its name on this socket. The name registered must have as its type part the string "AFPServer". [Note that server name string comparison is case insensitive].

Now a workstation trying to discover this or all file servers in the AppleTalk zone of interest must use NBP to lookup on the name =:AFPServer@<zone's name>. As a result of this call, the workstation client will receive, from NBP, a list of all active AFPServers in the zone together with the full internet address of the corresponding SLSs.

With this, the workstation has completed the discovery step of the login procedure. Note that so far there has been no need to open a session with any file server.

Obtaining File Server Information

After the SLS address has been obtained in the previous step, the AFP client in the workstation issues an FGetSvrInfo call to obtain a variety of server information including the list of AFPVersion strings and list of UAM strings for the versions and user access methods that the server can handle.

The Login Step

From the AFPVersion strings and UAM strings obtained in the previous step, the AFP client at this stage selects one AFPVersion string and one UAM string. These correspond to the AFP version and user authentication method the user wants the server to use. Then the user makes an FLogin call to AFP to initiate the login step. The caller must provide the selected AFPVersion string and UAM string, and of course the internet address of the SLS of the desired file server. Further information may have to be provided depending on the user authentication method selected.

The FLogin call causes the opening of an ASP session with the server and then an authentication of the user. If this proceeds to successful completion, then an AFP session is open with the server and further AFP calls can be sent to the server over this session.

The server associates with the user a 32-bit User ID and one or more 32-bit Group IDs, indicating the user's membership in those groups. In addition, one Group ID may be specially marked as the user's Primary Group, to be described below.

Chapter 4

Access Control Mechanisms in AFP

Access controls are built into AFP in three ways: user authentication at server login, an optional volume level password when first accessing ("opening") a volume, and directory-level access controls based on user authentication at login.

User Authentication at Server Login

User authentication at server login time has already been discussed above. This is the first "line of defense" as far as access controls are concerned. A file server can be set up to turn away an unauthorized user at this level and thus maintain the privacy and security of the file server's volumes and their contents. It is important to note that the user authentication step is vital to the directory-level access controls discussed below.

Volume Passwords

A second level of access control is provided through volume passwords. A server can associate an optional fixed-length 8-character password with each volume it is making visible through AFP.

As will become clear in the discussion of AFP calls, in order to make reference to a server volume the workstation must use a VolID. This VolID is obtained by the workstation through an FPOpenVol call in which the name of the volume is supplied as a call parameter. If the volume has a password associated with it, then the workstation must supply this password together with the FPOpenVol call in order to obtain the VolID (and hence the ability to issue AFP calls that make reference to that volume).

Volume passwords constitute a very simple protection mechanism for simple servers that do not wish to implement the full-fledged directory-level access control mechanism. Of course it is not as secure a mechanism either.

Directory-Level Access Controls

AFP includes a directory-level access control mechanism that constitutes the most secure method provided in this protocol. Note that, as mentioned earlier, AFP does not support any file-level access controls.

The basic idea is that directories are containers that contain objects: other containers (directories) and files. A user that has been authenticated at the server login step can try to make four classes of accesses to the directory's contents:

- *see directories* contained within that directory
- *see files* contained within that directory

- *read the contents of an object* (or view its parameters) contained within that directory.
- *write the contents of an object* (or change its parameters) contained within that directory.

Permission to search a directory's contents allows the user to list the names and parameters of other files and/or directories contained in this directory. Access to the contents of an objects in a directory can be of two types: Read and/or Write. Setting (changing) the parameters of an object (file or directory) is considered as Writing to the object.

AFP includes mechanisms for restricting these classes of access (Search, Read, and Write access) at the level of each directory in the catalog tree. This is done as follows:

With each directory we associate an owner and a group of users. The owner is initially set to the ID of the user that created the directory; the group is initially set to the user's Primary group, if one exists. The Primary group is like any other group affiliation, except that it is the group that gets assigned to any new directories that the user creates.

AFP recognizes three different access rights: Search, Read, and Write. The file server must save, with each directory, three *Access Rights bytes*, corresponding respectively to the directory's owner, the directory's group, and the world. Each access rights byte is used as a bit map whose three least significant bits are used to encode the three different access rights allowed by AFP for the corresponding user (i.e. owner, group, or world). More specifically, each directory has associated with it the following five parameters:

- *owner ID* [4 bytes]
- *group ID* [4 bytes]
- *owner's Access Rights* [1 byte]
- *group's Access Rights* [1 byte]
- *world's Access Rights* [1 byte].

In addition, the server must maintain a one-to-one mapping between owner ID (user ID) and user name (a string of maximum length 31 characters), and between group ID and group name (also a string of at most 31 characters). AFP includes calls to allow users to map IDs to names and visa-versa.

The most significant 5 bits of each access rights byte must be zero (they are reserved for access rights extension in future versions of this protocol).

When a user logs in on a server, as part of the user authentication mechanism, various identifiers are retrieved from a *user data base* maintained on the server. These include the user ID (a 32-bit number unique among all server users) and one or more (the exact number is implementation dependent) 32-bit group IDs which indicate the user's group affiliations. One of these group IDs is special in that it represents the user's Primary Group. This number is assigned as the group ID of each directory created by the user. Assignment of user IDs, group IDs, and Primary Groups is an administrative function and is outside the scope of this document.

When this user tries to access a directory or its contents, the following algorithm is used by the server to extract the rights corresponding to that user (UARights) and that directory:

```
UARights := Directory's world access rights;
If (User's ID = Directory's owner ID) then
```

UARights := UARights OR Directory's owner's access rights;
 If (any of User's group IDs = Directory's group ID) then
 UARights := UARights OR Directory's group's access rights

The OR operations in this algorithm are inclusive OR operations. Having given this algorithm for determining the UARights corresponding to a particular directory we can now examine in more detail what rights are required for various AFP operations. We use the following notation: **SA** = "search access rights to all ancestors down to but not including the Parent directory", **WA** = "search or write access to all ancestors down to but not including the Parent directory", **Sp** = "search access rights to the Parent directory", **Wp** = "write access rights to the Parent directory", **Rp** = "read access rights to the Parent directory".

Almost all operations require **SA** access, which means the user can only access the contents of an object in a given directory if he has permission to Search every directory in the path from the root to the parent's immediate parent directory. The exact access permitted to objects in the directory is then determined further by **Sp**, **Rp** and **Wp** rights for the Parent directory.

Creating a File or a Directory: The user must have **WA** plus **Wp**. Hard Create needs the same rights as deleting a file.

Enumerate a Directory: The user must have Search access rights to all directories down to but not necessarily including the directory being enumerated. To view its offspring that are directories, Search access to the directory being enumerated is required as well. To view its file offspring, Search access to the directory is not required, but the user must have Read access rights to the directory.

Deleting a File. The user must have **SA**, **Rp**, and **Wp**. Furthermore, a file can be deleted only if it is not open at that time.

Delete a Directory: The user must have **SA**, **Sp**, and **Wp**. Furthermore, a directory can be deleted only if it is empty.

Rename a File: The user must have **SA**, **Rp**, and **Wp**.

Rename a Directory. The user must have **SA**, **Sp**, and **Wp**.

Get (Read) Directory Parameters: All this requires are **SA** plus **Sp** access rights.

Get (Read) File Parameters: The user must have **SA** plus **Rp**.

Open a File to Read its Contents: A file's fork must be opened in read mode before its contents may be read. To open a file for read (and thus to be able to read from it) the user must have **SA** plus **Rp**.

Open a File to Write its Contents: A file's fork must be opened in write mode in order to write to the fork. To open a fork which is currently empty (both forks must be of zero-length) to write to it, the user must have **WA** plus **Wp**. To open an existing fork (either fork is of non-zero length) to write to it, **SA**, **Rp**, and **Wp** access is required.

Set (Write) File Parameters: The user must have **WA** plus **Wp** to set the parameters of an empty file (when both forks are zero-length). To set the File Parameters on a non-empty (either fork) file, **SA**, **Rp**, and **Wp** access is required.

Set (Write) Directory Parameters: The user must have **SA**, **Sp**, and **Wp** access to change a directory's parameters if the directory is non-empty. If the directory is empty, the user must have **WA** plus **Wp** to change its parameters.

Move a Directory or a File: Through AFP a directory or a file can be moved from a source parent directory to a destination parent directory on the same volume. The user must have **SA** rights to the source parent directory, **WA** access to destination parent directory, plus **Write** access rights to both source and destination parents. Furthermore to move a file, the user needs in addition **Read** access rights to the source parent directory. To move a folder, **Search** access to the source parent directory is required in addition instead of **Rp**.

Modify a Directory's Access Rights Information: A directory's Owner ID, Group ID, and three Access Rights bytes can be modified only if the user is the directory's owner and then only if the user has **WA** plus **Wp** or **Sp** access to the parent directory.

Copy a file (FPCopyFile): To copy a file, possibly across volumes managed by the server, the user must have **SA** plus **Rp** to the source parent directory and **WA** plus **Wp** to the destination parent directory.

Special Cases

OwnerID=0 means that the folder is "unowned" or owned by <any user>. The owner bit of the User Rights summary byte is always set for such a folder.

GroupID=0 means that the folder has no group affiliation; hence the group's access privileges (R, W, S) are ignored for such a folder.

Chapter 5

A Discussion of AFP Calls

Now we provide an overall discussion of the various calls provided by AFP and how they can be used to access a file server. A completely detailed specification of each call is available in Chapter 7. For the purpose of this discussion we classify the calls into various groups.

Server-Level Calls

A workstation client can use the following server-level AFP calls:

- *FPGetSrvrInfo*
- *FPLogin*
- *FPLoginCont*
- *FPGetSrvrParms*
- *FPLogout*
- *FPMapID*
- *FPMapName*.

To begin, a workstation uses the Name Binding Protocol to discover the server's session listening socket's network address (we call this the *SAddr*).

Next, the workstation must obtain server information by using the *FPGetSrvrInfo* call. This is done without opening a session to the server and it returns a block of server information containing the following server parameters:

- the server's name
- a string identifying the server's machine type
- a list of the AFP versions the file server can handle
- a list of the various user authentication methods the server can handle
- an icon to be used for displaying server volumes on the Macintosh.
- a bitmap of flags.

After making this call the workstation's AFP client selects one AFP version and one user authentication method. Using this selection it makes the *FPLogin* call to establish a session with the file server. A session is needed before any of the other calls can be made to the server. This login step involves authentication of the user; it returns an *SRefNum*, a session reference number to be used in all calls made over this session. Depending on the chosen authentication method, the entire authentication process may involve additional *FPLoginCont* (Login continue) calls to provide more information to the server.

After a session has been established with the server, the workstation must obtain a list of the volumes on the server. This is done by making the *FPGetSrvrParms* call, which returns a count of the number of such volumes, the names of these volumes, and an indication of whether or not these volumes are password-protected.

When a workstation no longer needs to communicate with a server it issues an *FPLogout* call to terminate the session.

The *FMapID* and *FMapName* calls relate to access controls issues. The *FMapID* call is used to obtain the User or Group Name corresponding to a given User or Group ID. The *FMapName* call provides the inverse functionality, mapping a User or Group Name into the corresponding User or Group ID.

Volume-Level Calls

There are five volume-level AFP calls:

- *FPOpenVol*
- *FPCloseVol*
- *FPGetVolParms*
- *FPSetVolParms*
- *FPFlush*.

After obtaining the volume names through the *FPGetSvrParms* call, the workstation client must make an *FPOpenVol* call for each volume it wishes to have access to. If the volume has a password attached to it, it must be supplied at this time. The call returns the volume parameters asked for in the call, the key one being the *VolID*. This *VolID* is used in all subsequent calls to identify the volume to which those calls apply.

The *VolID* remains a valid identifier either until the session is terminated or until an explicit *FPCloseVol* call is made. This call invalidates the *VolID*.

After obtaining a volume's *VolID* the workstation client can, at any time, obtain the volume's parameters by making an *FPGetVolParms* call. Likewise, the volume's parameters may be changed through an *FPSetVolParms* call.

A request may be made to the server to write out to its disk any data and control structures pertinent to a particular volume. This is done by making an *FPFlush* call.

Directory-Level Calls

There are five directory-level AFP calls:

- *FPSetDirParms*
- *FPOpenDir*
- *FPCloseDir*
- *FPEnumerate*
- *FPCreateDir*.

The *FPSetDirParms* call allows the workstation client to modify a directory's parameters. The *FPGetFileDirParms* call (discussed below) is used to obtain a directory's parameters from the file server.

The *FPOpenDir* call is used to "open" a directory on a Variable-DirID volume and retrieve its *DirID*, which can then be used in subsequent calls to enumerate the directory or access its offspring. For Variable-DirID volumes, this is the only way to retrieve the *DirID* (using a *FPGetFileDirParms* or *FPEnumerate* call to retrieve the *DirID* on such volumes will return an error).

Note that it is not considered an error to make this call for directories on Fixed-DirID volumes (the fixed DirID will be returned), but this is not the preferred method. Use *FPGetFileDirParms* or *FPEnumerate* instead. Directories on variable DirID volumes can be "closed" by making an *FPCloseDir* call, which invalidates the corresponding DirID.

The *FPEnumerate* call is one of the most important AFP calls. It is used to enumerate (i.e., list) the objects (files and directories) contained within a specified directory. In reply to this call the server returns a list of directory and/or file parameters corresponding to these objects.

Directories are created through the *FPCreateDir* call.

Each of these calls requires the specification of the directory to which the call applies. This specification, in general, consists of the VolID, a DirID (of the directory or of an ancestor), and (if this DirID is of an ancestor, then) a pathname down to the specified directory. Remember that the DirID of the root of a volume is always equal to 2.

File-Level Calls

There are three file-level calls:

- *FPSetFileParms*
- *FPCreateFile*
- *FPCopyFile*.

The *FPSetFileParms* call is used to modify a specified file's parameters. The *FPGetFileDirParms* call (discussed below) is used to obtain a specified file's parameters.

A file can be created through the *FPCreateFile* call.

A file that exists on a volume managed by a server can be copied to any other volume managed by that server with the *FPCopyFile* call.

Combined Directory-File-Level Calls

AFP includes five calls that apply to files and directories:

- *FPGetFileDirParms*
- *FPSetFileDirParms*
- *FPRename*
- *FPDelete*
- *FPMove*.

The *FPGetFileDirParms* call is used to retrieve the parameters associated with a given object. Using this call, one need not know in advance whether the object is a file or a directory; indication of its type is returned from the call. Likewise, the *FPSetFileDirParms* call is used to set certain parameters associated with a given object, even if one does not know in advance whether the object is a file or directory. This call only allows the setting of those parameters that are common to both types of object.

The next two calls can be used to respectively rename or delete files and directories. A file can be deleted only if it is not open; a directory can be deleted only if it is empty.

The *FPMove* call can be used to move a file or a directory from one parent directory to another on the same volume. At the same time the object moved can be renamed.

Fork-Level Calls

There are eight fork-level calls:

- *FPGetForkParms*
- *FPSetForkParms*
- *FPOpenFork*
- *FPCloseFork*
- *FPRead*
- *FPWrite*
- *FPFlushFork*
- *FPByteRangeLock*.

A fork's parameters can be read or modified by using the *FPGetForkParms* and *FPSetForkParms* calls.

The *FPOpenFork* call is used to open an existing file's fork (either one). This call returns an *OForkRefNum* which is used in all calls that refer to this open fork. This reference number stays valid until the fork is closed through the *FPCloseFork* call.

The four remaining calls can be used to manipulate a fork opened through a previously issued *FPOpenFork* call. The contents of the fork can be read by making *FPRead* calls. The client can write to a fork by using *FPWrite* calls. The *FPFlushFork* call makes the server flush (write out to its disk) any of that fork's data that is in the server's internal buffers.

To allow for shared use of a file's open fork, a client can lock ranges of bytes in that fork. If a client locks a particular byte range (through the *FPByteRangeLock* call) then that range of bytes is reserved for exclusive manipulation by the client placing the lock; other clients can neither read nor write within the locked range.

Chapter 6

A Design for The Finder's Desktop Management In a Network Environment

The Finder presents Macintosh user with a unique user interface centered around the use of icons to represent objects on a disk volume. To present this interface to the user the Finder makes use of a number of data structures separate from the File System's volume catalog, all of which are maintained as resources of various types in an invisible resource file called 'Desktop'.

The Desktop file is currently used to perform three separate functions:

1. To associate documents and applications with particular icons through its 'bundle' mechanism, as well as storing the actual icon bitmaps,
2. To locate the corresponding application when a user opens a document, and
3. To store the text of comments associated with files and directories as part of the information displayed by 'Get Info'.

The management of icons centers around the concept of a *bundle*, stored as a resource of type BNDL in an application's resource fork. The BNDL resource, which is identified with a particular FileCreator type can, among other things, refer to a number of FREF resources, which can in turn indicate that a file of a particular type should be displayed using a particular icon. Together, BNDLs and FREFs can be used to determine the icon to be displayed for a given file from the Creator and Type information stored as part of its Finder Info in the catalog.

In addition, the Desktop file is used on HFS volumes to hold a list of applications stored in subdirectories on the volume. The desktop contains an APPL resource which is used by the Finder to find an application to launch when a document is selected, given the document's FileCreator information. The APPL resource basically maps a particular Creator type to a list of applications that can open documents of the specified type.

Finally, the Desktop file is used as a repository for the text of comments associated with files and directories on the volume. Comments are retrieved when the user selects 'Get Info' for a file or directory, at which time the comment text can also be changed.

The use of the current Finder in a Network Environment (i.e. on File Server volumes) has proven unsatisfactory: resource files, such as the Finder's Desktop file, are ill-suited for sharing among multiple users on a single File Server.

A new mechanism has been designed to replace the Finder's direct use of the Desktop resource file on File Server volumes completely. Eventually, this mechanism could be used transparently for both local and remote volumes. The interface to this new Desktop mechanism is presented below.

Call Interface

The interface includes three groups of calls:

1. Icon calls (AddIcon, GetIcon, and GetIconInfo),
2. APPL calls (AddAPPL, DeleteAPPL, and GetAPPL),
3. Comment calls (AddComment, DeleteComment, GetComment).

Each call is mapped into a corresponding AFP command. The semantics of the AFP calls exactly parallel the semantics of the interface routines. Like FPWrite, FPAAddIcon is special in that it requires a special intermediate exchange of packets to transfer the data block representing the icon bitmap.

In the descriptions of the individual calls in the following sections, the data type **ResType** refers to the 4-byte signatures that are part of every file's Finder Information. Each file is assigned a 'File Type' meant to be representative of the nature of the contents of the file (PNTG, TEXT, etc.) and a 'File Creator', which is a unique signature indicating the application which created the file (such as MPNT, MACA, etc.).

Before any Desktop calls can be made, the user must make an OpenDT call, as follows:

- **Function** OpenDT(VolID: Integer; Var DTRefNum: Integer): OSErr;

The file RefNum returned for the Desktop Database must be used on future calls to indicate the Desktop Database being referred to. If an error occurred on the call, the refNum returned will be zero.

When all Desktop operations have been completed, the user should make a CloseDT call (which takes a single argument, the DTRefNum) and returns an OSErr. This will free all resources allocated as part of the OpenDT call.

Icon Related Calls

- **Function** AddIcon(DTRefNum: Integer; FileCreator, FileType: ResType; IconType: Byte; IconTag: LongInt; BitmapSize: Integer; Bitmap: Ptr): OSErr;

AddIcon adds a new icon bitmap to the Desktop database. The FileType and FileCreator arguments (4 bytes each) specify the set of files this icon is associated with, while the IconType argument may indicate a specific kind of icon. Note that for a given FileCreator/FileType, there may be a number of icons available, each with a different IconType. The IconTag argument indicates a LongInt value to be associated with the icon which will be returned along with the icon bitmap when it is retrieved. This could be used as a timestamp, for instance, to associate the creation date of the application with the icons it exports. Finally, the Size and Bitmap arguments provide the actual bitmap in questions.

If an icon of the specified IconType already exists for the indicated FileCreator and FileType, **AddIcon** will replace the bitmap stored with the new Bitmap. An error will be returned if the size of the new bitmap is different from the size of the old bitmap.

- **Function** GetIcon(DTRefNum: Integer; FileCreator: ResType; FileType: ResType; IconType: Byte; Var Length: Integer; BitMap: Ptr): OSErr;

GetIcon retrieves the bitmap for a given icon, given its FileCreator, FileType and IconType. If an icon of type IconType of the specified FileCreator and FileType is available, it is returned. Otherwise, an ItemNotFound error is returned. The length argument used on input to indicate the size of the buffer pointed to by the BitMap pointer. When the call is completed it is overwritten with the actual size of the bitmap returned.

- **Function GetIconInfo(DTRefNum: Integer; FileCreator: ResType; IconIndex: Integer; Var IconTag: LongInt; Var FileType: ResType; Var IconType: Byte; Var Size: Integer): OSErr;**

GetIconInfo retrieves a description of an icon, given its FileCreator type and a numerical index. It can be used to determine the set of icons associated with a given application without knowing the FileTypes in advance. Successive calls with increasing values of IconIndex will return information on all icons associated with a given Creator type.

Application Related Calls

- **Function AddAPPL(DTRefNum: Integer; FileCreator: ResType; DirID: LongInt; CName: String[31]; APPLTag: LongInt): OSErr;**

AddAPPL adds an entry for the application specified by the DirID/CName under the indicated ResType. The APPLTag argument is an additional LongInt stored with the mapping information. There may be more than one application with same FileCreator ResType, although the DirID/CName should uniquely identify the file. The Tag information might be used to decide among many possible applications which one to launch for a particular document (if the tag of the creator were stored in the Finder information of the document, for instance).

- **Function RemoveAPPL(DTRefNum: Integer; FileCreator: ResType; DirID: LongInt; CName: String[31]): OSErr;**

RemoveAPPL removes the mapping information for a given application indicated by its DirID/CName. Note that while the FileCreator type must be specified to locate the entry, the application tag is not required to remove an application entry.

Note that it is the Finder's responsibility to add and remove entries for applications which are copied to the volume or deleted, respectively. For entries which are moved or renamed, the Finder should remove the entry before the operation and add a new entry with the updated information after the operation has been completed successfully.

- **Function GetAPPL(DTRefNum: Integer; FileCreator: ResType; Index: Integer; Var APPLTag: LongInt; Var DirID: LongInt; Var CName: StringPtr): OSErr;**

GetAPPL looks up an application given its Creator ResType. The index argument is used to enumerate all application mappings stored. Indices 1 through n will retrieve the 1st through nth application mapping stored which are accessible by the caller (i.e. to which the user has Search and Read access). Unless the caller wishes to implement a special selection algorithm over all available applications, a single call to get the first mapping should suffice to find an application which can be launched to open the selected document.

Comment Related Calls

- **Procedure AddComment**(DTRefNum: Integer; DirID: LongInt; CName: String[31]; CommentText: String[199]);

AddComment stores a comment string associated with a particular file or directory on the volume. Unlike icons, there can be no more than one comment associated with any file or directory. If **AddComment** is called for a file or directory which already has an associated comment, the existing comment is replaced.

- **Function RemoveComment**(DTRefNum: Integer; DirID: LongInt; CName: String[31]): OSErr;

RemoveComment removes the comment associated with a particular file or directory. An error is returned if no comment was stored for the file or directory.

Note that while the Finder will call **RemoveComment** to remove comments for files or directories when they are deleted, it does not call **GetComment**, **RemoveComment** and **AddComment** whenever a file or directory is renamed or moved.

- **Function GetComment**(DTRefNum: Integer; DirID: LongInt; CName: String[31]; Var CommentText: String[199]): OSErr;

GetComment retrieves the comment associated with a particular file or directory. If a comment is stored, the comment text is returned. If no comment exists, an error is returned.

Chapter 7

Specification of AFP Calls

This section provides a specification of the various AFP calls.

For each call a brief description of the call is provided. This is followed by a list of the input and output parameters of the call. For each call, the underlying transport mechanism is assumed to return an *FPErr* whose relevant values for the particular call are listed and explained. For all calls, if the call completes successfully, then the error code value *NoErr* is returned. This particular value of *FPErr* is not included in the descriptions of the calls. Likewise, a *UserNotAuth* error can be returned from almost every call, but is not included in the call descriptions. This error indicates that the user has not yet been properly logged in. In addition, a *MiscErr* can be returned from almost every call. *MiscErr* is used by the server to map other errors that don't have an equivalent AFP Error (like error in reading a disk sector).

A description of the relevant algorithm pertaining to the call is also included with the discussion. The access rights required by each call are also specified.

Each AFP call is sent to the server in the form of a *Command* block, to which the server responds with the four-byte *FPErr* plus a *Reply* block. For each call, the formats of the Command and Reply blocks are provided in pictorial form. The pictures are drawn so that the width of the rectangles is one byte. The bit numbering and byte ordering of multi-byte fields is performed according to MC68000 standards: in the following pictures, high byte always appears above low byte. How these blocks translate into network packets depends on the underlying transport mechanism used by AFP. It should be noted that some calls return an empty Reply block (in these cases this block is not shown).

The *FPWrite* call is special in that the data to be written is not included in the Command block but is expected to be passed to the underlying transport mechanism as a separate block.

Many of the calls require a bitmap to be passed, along with a block of parameters packed in "Bitmap order." This order is defined as follows: the parameter corresponding to the least-significant set bit in the bitmap is packed first, followed by the parameter corresponding to the next-more-significant bit, etc., ending with the parameter corresponding to the most-significant bit.

All numerical fields represent signed numbers unless otherwise indicated.

FPAddAPPL

This call adds an APPL mapping to the Desktop Database.

INPUTS:

SRefNum (INT)	session refnum
DTRefNum (INT)	Desktop Database refNum
Directory ID (LONG)	ancestor directory identifier
FileCreator (RESTYPE)	creator type of application being added
APPLTag (LONG)	a user-defined tag stored with the APPL entry
PathType (BYTE)	indicates whether pathname is composed of long names or short names: 1 = all pathname elements are short names 2 = all pathname elements are long names
PathName (STRING)	pathname to the application being added

OUTPUTS: FPErrors (LONG)


ERRORS:

ObjectNotFound	input parameters do not point to an existing file
AccessDenied	user does not have the rights listed below.

ALGORITHM: A mapping to the specified application is added to the Desktop Database. If an entry for the same application (same file name in the same directory with the same file creator) already exists, it is replaced.

RIGHTS: The user must have previously called FPOpenDT for the corresponding volume. In addition, the application must be present in the specified directory before this call is issued, and the user must have Search or Write access rights to all ancestors except the object's Parent, as well as Write access rights to the Parent.

PACKET FORMAT:



FPAddComment

This call adds a comment for a file or directory to the Desktop Database.

INPUTS:	SRefNum (INT)	session refnum
	DTRefNum (INT)	Desktop Database refnum
	Directory ID (LONG)	directory identifier
	PathType (BYTE)	indicates whether pathname is composed of long names or short names: 1 = all pathname elements are short names 2 = all pathname elements are long names
	PathName (STRING)	pathname to the file or directory with comment
	CmtLength (BYTE)	length of comment data
	CmtText (BYTES)	comment data to be associated with file or folder specified (limited to 199 bytes)

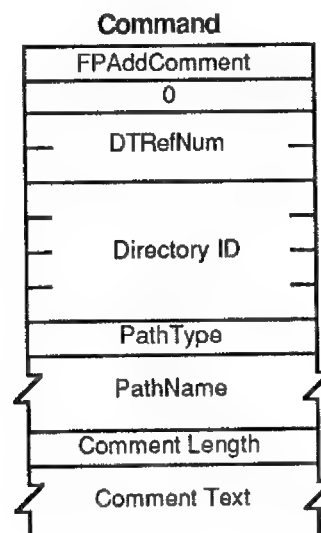
OUTPUTS: FPErr (LONG)

ERRORS:	ParamErr	unknown session refnum or Desktop Database refnum; bad pathname
	ObjectNotFound	input parameters do not point to an existing file or dir
	AccessDenied	user does not have the rights listed below

ALGORITHM: The comment type and comment data are associated with the specified file or directory and stored in the Desktop Database. If the comment length is greater than 199 bytes, the comment will be truncated to 199 bytes and no error will be returned.

RIGHTS: The user must have previously called FPOpenDT for the corresponding volume. In addition, the object must be present in the specified directory before this call is issued, and the user must have Search or Write access rights to all ancestors except the object's Parent, as well as Write access to the Parent.

PACKET FORMAT:



FPAddIcon

This call is used to add an icon bitmap to the Desktop Database.

INPUTS:	SRefNum (INT)	session refnum
	DTRefNum (INT)	Desktop Database refnum
	FileCreator (RESTYPE)	file's creator type
	FileType (RESTYPE)	file's type
	IconType (BYTE)	type of icon being added
	IconTag (LONG)	tag information to be stored with the icon
OUTPUTS:	BitmapSize (INT)	size of the bitmap for this icon
	FPErrror (LONG)	
ERRORS:	ParamErr	unknown session refnum or Desktop Database refnum
	IconTypeError	new icon size is different from existing icon's size
	AccessDenied	user does not have the rights listed below
ALGORITHM:	A new icon is added to the Desktop database for the specified FileCreator and FileType. If an icon of the same FileCreator, FileType, and IconType already exists the icon is replaced. If the new icon's size is different from the old icon's size an IconTypeError is returned.	
RIGHTS:	The user must have previously called FPOpenDT for the corresponding volume. In addition, the volume indicated by DTRefnum must not be marked ReadOnly.	
NOTES:	The command packet includes all input parameters except for the actual bitmap. The bitmap is sent to the server in a subsequent intermediate exchange of Session Protocol packets.	
PACKET FORMAT:		

Command	
FPAAddIcon	
0	
DTRetNum	
FileCreator	
FileType	
IconType	
0	
IconTag	
IconSize	

FPByteRangeLock

This call is used to lock a range of an open fork to ensure exclusive access. Locks prevent all other users from reading or writing any bytes within the range.

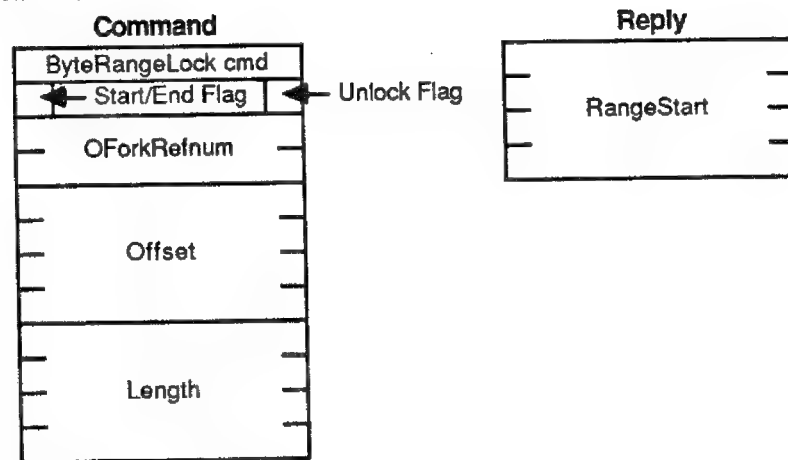
INPUTS:	SRefNum (INT)	session refnum
	OForkRefnum (INT)	open fork refnum
	Offset (LONG)	offset to the first byte of the range to be locked or unlocked (can be negative if Start/End Flag = End)
	Length (LONG)	number of bytes to be locked or unlocked (signed—can't be negative)
	UnlockFlag (BIT)	flag to indicate whether range is to be locked or unlocked: 0 = lock 1 = unlock
	Start/EndFlag (BIT)	flag indicating whether the Offset field is relative to the beginning or end of the fork (valid only when locking) - (all other bits must be zero): 0 = relative to the beginning of the fork 1 = relative to the end of the fork
OUTPUTS:	FPErr (LONG)	
	RangeStart (LONG)	number of the first byte of the range just locked (valid only when returned from a successful lock command)
ERRORS:	ParamErr	unknown session refnum or open fork refnum; combination of Start/End flag and offset specified a range starting before the 0th byte
	LockErr	some or all of requested range is locked by another user
	NoMoreLocks	server's maximum lock limit has been reached
	RangeOverlap	user attempted to lock (some or all of) a range that is already locked by the user
	RangeNotLocked	tried to unlock a range that was not locked by the user

ALGORITHM: If no other user holds a lock on any part of the requested range, the server will lock exactly the specified range for this user. A user may hold multiple locks on a given open fork, up to a server-specific limit. Multiple locks may not overlap. A lock range may start and/or extend past the end-of-fork; this does not prevent another user from writing to the fork past the locked range. Specifying an Offset of zero and a Length of \$FFFFFFFF will lock the entire fork. All locks held by a user are unlocked when the user closes the fork. Unlocking a range makes it available for reading and writing to other users. A RangeNotLocked error is returned if an attempt was made to unlock a range that was not locked by the user.

If multiple writers are concurrently modifying the fork, they may each have a different notion of the end-of-fork, although the server always knows the correct end-of-fork. It is for this reason that the "lock relative to end-of-fork" feature is provided. The number of the first locked byte is returned, since the end-of-fork may have been different from the user's notion at the time at which the call was made.

RIGHTS: No special access rights are needed to make this call.

PACKET FORMAT:



FPCloseDir

This call is used to close a directory.

INPUTS: SRefNum (INT) session refnum
VolumeID (INT) volume identifier
DirID (LONG) ancestor directory identifier

OUTPUTS: FPErr (LONG)

ERRORS: ParamErr unknown session refnum or volume identifier
ObjectNotFound unknown directory identifier

ALGORITHM: The directory identifier is invalidated, and may not be used again.

RIGHTS: The user must have previously called FPOpenVol for this volume and FPOpenDir for this directory.

PACKET FORMAT:

Command	
CloseDir command	
Ø	
Volume ID	
Directory ID	

FPCloseDT

This call is used to disassociate a user from the volume's Desktop Database.

INPUTS: SRefNum (INT) session refnum
DTRefNum (INT) Desktop Database refNum, as returned from FPOpenDT

OUTPUTS: FPErrror (LONG)

ERRORS: ParamErr unknown session refnum or Desktop Database refnum

RIGHTS: The user must have made a successful FPOpenDT call before this call can be made.

PACKET FORMAT:

Command	
FPCloseDT	
0	
—	DTRefNum —

FPCloseFork

This call is used to close a fork which was opened by FPOpenFork.

INPUTS: SRefNum (INT) session refnum
OForkRefnum (INT) open fork refnum

OUTPUTS: FPErrors (LONG)

ERRORS: ParamErr unknown session refnum or open fork refnum

ALGORITHM: The server flushes and then closes the open fork, invalidating the OForkRefnum. If the fork had been written to, the file's Mod Date will be set to the server's clock at this time.

RIGHTS: No special access rights are needed to make this call.

PACKET FORMAT:

Command	
CloseFork command	
Ø	
—	OForkRefnum —

FPCloseVol

This call is used to "unmount" a volume.

INPUTS: SRefNum (INT) session refnum
VolumeID (INT) volume identifier

OUTPUTS: FPErrror (LONG)

ERRORS: ParamErr unknown session refnum or volume identifier

ALGORITHM: The Volume ID is invalidated. No further calls may be made to access objects on this volume unless another FPOpenVol call is made.

RIGHTS: The user must have previously called FPOpenVol for this volume.

PACKET FORMAT:

Command	
CloseVol command	
Ø	
—	Volume ID —

FPCopyFile (optional; may not be supported by all servers)

This call is used to copy a file residing on one of the server's volumes to another location on one of the server's volumes. The destination of the copy is specified by providing a VolID, DirID, and Pathname that indicate the copy's new Parent Directory.


INPUTS:	SRefNum (INT)	session refnum
	SVolumeID (INT)	source volume identifier
	SDirID (LONG)	source ancestor directory identifier
	SPathType (BYTE)	indicates whether SPathname is composed of long names or short names: 1 = all pathname elements are short names 2 = all pathname elements are long names
	SPathname (STR)	pathname of the file to be copied (cannot be null)
	DVolumeID (INT)	destination volume identifier
	DDirID (LONG)	destination ancestor directory identifier
	DPathType (BYTE)	indicates whether DPathname is composed of long names or short names (same values as SPathType)
	DPathname (STR)	pathname to the destination Parent Directory (may be null)
	NewType (BYTE)	indicates whether NewName is a long name or a short name (same values as SPathType)
	NewName (STR)	name to be given to the copy (may be null)
OUTPUTS:	FPErr (LONG)	
ERRORS:	ParamErr	unknown session refnum, volume identifier or pathname type; bad pathname or NewName
	ObjectNotFound	the source file does not exist; unknown ancestor directory
	ObjectExists	an object by the name NewName already exists in the destination Parent Directory
	AccessDenied	user does not have the right to read the file or write to the destination
	CallNotSupported	call not supported by this server
	DenyConflict	the file cannot be opened for Read, Deny Write
	DiskFull	no more space on the volume
	ObjectTypeErr	source parameters point to a directory

ALGORITHM: The server will attempt to open the source file for Read, Deny Write. If that fails, a Deny Conflict error will be returned. Otherwise, the file will be copied from source to destination. The original file is not changed or deleted.

The copy is given the name specified in NewName. If NewName is null, the server will attempt to give the copy the same name as the original. The creation of Long and Short names is performed as described in Appendix B. A unique FileNumber is assigned to the file. Its Parent Directory ID will be set to the Dir ID of the destination Parent Directory. All other file parameters remain the same as the source file's parameters. The Mod Date of the destination Parent Directory is set to the server's clock.

RIGHTS: The user must have previously called FPOpenVol for both source and destination volumes. In addition, the user must have Search access rights to all ancestors except the source file's Parent Directory, and Read access right to the source file's Parent directory. Further, the user must have Search access rights to all ancestors except the destination Parent Directory, and Write access right to the destination Parent Directory.

PACKET FORMAT:



FPCreateDir

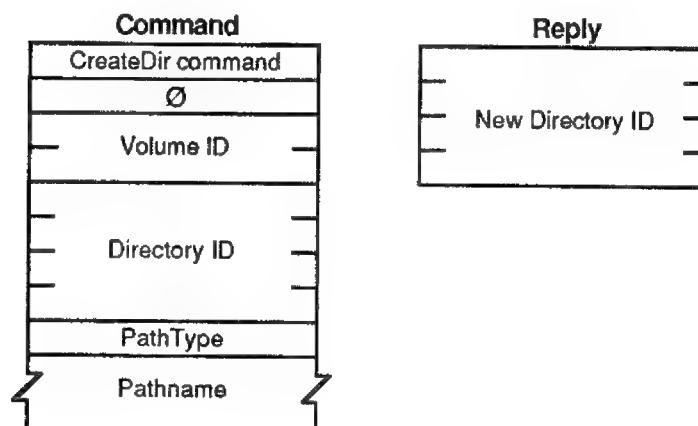
This call is used to create a new directory.

INPUTS:	SRefNum (INT)	session refnum
	VolumeID (INT)	volume identifier
	DirID (LONG)	ancestor directory identifier
	PathType (BYTE)	indicates whether Pathname is composed of long names or short names: 1 = all pathname elements are short names 2 = all pathname elements are long names
	Pathname (STR)	pathname to desired directory (cannot be null)
OUTPUTS:	FPErr (LONG)	
	NewDirID (LONG)	identifier of new directory
ERRORS:	ParamErr	unknown session refnum, volume identifier, or pathname type; null or bad pathname
	ObjectNotFound	unknown ancestor directory
	ObjectExists	an object already exists by that name
	AccessDenied	user does not have the rights listed below
	FlatVol	the volume is flat and does not support directories
	DiskFull	no more space on the volume

ALGORITHM: If the volume is not flat, a new empty directory is created with the name as specified in Pathname. A unique NewDirID is assigned to the directory. Its OwnerID is set to the UserID of the user making the call, and its GroupID is set to the ID of the user's Primary Group. Access rights are initially set to Read, Write, and Search for the Owner, no rights for the Group or World. Finder Info is zeroed. Create Date and Mod Date are set to the server's clock. Backup Date is set to \$80000000, signifying that this directory has never been backed up. The creation of Long and Short names is performed as described in Appendix B. All attributes are initially cleared. The Mod Date of the Parent Directory is set to the server's clock.

RIGHTS: The user must have previously called FPOpenVol for this volume. In addition, the user must have Search or Write access rights to all ancestors except this directory's Parent Directory, as well as Write access right to the Parent Directory.

PACKET FORMAT:

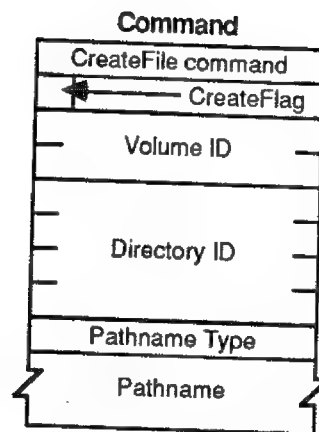


FPCreateFile

This call is used to create a file.

<i>INPUTS:</i>	SRefNum (INT)	session refnum
	VolumeID (INT)	volume identifier
	DirID (LONG)	ancestor directory identifier
	CreateFlag (BIT)	a flag that specifies hard or soft create (all other bits will be zero): 0 = soft create 1 = hard create
	PathType (BYTE)	indicates whether Pathname is composed of long names or short names: 1 = all pathname elements are short names 2 = all pathname elements are long names
	Pathname (STR)	pathname including name of new file (cannot be null)
<i>OUTPUTS:</i>	FPErr (LONG)	
<i>ERRORS:</i>	ParamErr	unknown session refnum, volume identifier, or pathname type; null or bad pathname
	ObjectNotFound	unknown ancestor directory
	ObjectExists	soft create: a file by that name already exists
	ObjectTypeErr	a directory by that name already exists
	AccessDenied	user does not have the rights listed below
	FileBusy DiskFull	hard create: the file already exists and is open no more space on the volume
<i>ALGORITHM:</i>	In a soft create, if the object does not already exist, a new file is created with the name as specified in Pathname. A unique FileNumber is assigned to the file. Finder Info is zeroed. Create Date and Mod Date are set to the server's clock. Backup Date is set to \$80000000, signifying that this file has never been backed up. The creation of Long and Short names is performed as described in Appendix B. The lengths of both forks are set to zero. The Mod Date of the file's Parent Directory is set to the server's clock. All file attributes are initially cleared. The Mod Date of the Parent Directory is set to the server's clock.	
	In a hard create, if the file already exists, it is essentially deleted and then recreated. All file parameters (including the Create Date) are reinitialized as described above.	
<i>RIGHTS:</i>	The user must have previously called FPOpenVol for this volume. For a soft create, the user must have Search or Write access rights to all ancestors except this file's Parent Directory, as well as Write access right to the Parent Directory. For a hard create, the user must have Search access to all ancestors except the Parent, as well as Read and Write access to the Parent.	

PACKET FORMAT:




FPDelete

This call is used to delete either a directory or file.

INPUTS:	SRefNum (INT)	session refnum
	VolumeID (INT)	volume identifier
	DirID (LONG)	ancestor directory identifier
	PathType (BYTE)	indicates whether Pathname is composed of long names or short names: 1 = all pathname elements are short names 2 = all pathname elements are long names
	Pathname (STR)	pathname of file or directory to be deleted (may be null if a directory is to be deleted)
OUTPUTS:	FPError (LONG)	
ERRORS:	ParamErr	unknown session refnum, volume identifier, or pathname type; bad pathname
	ObjectNotFound	input parameters do not point to an existing file or directory
	DirNotEmpty	the directory is not empty
	FileBusy	the file is open
	AccessDenied	user does not have the rights listed below
ALGORITHM:	If the object to be deleted is a directory, the server checks to see if it contains any offspring; a DirNotEmpty error is returned if so. If a file is to be deleted, it must not be currently open by any user (else a FileBusy error is returned). The Mod Date of the object's Parent Directory is set to the server's clock.	
RIGHTS:	The user must have previously called FPOpenVol for the volume. In addition, the user must have Search access rights to all ancestors except the object's Parent Directory, as well as Write access right to the Parent Directory. If a directory is being deleted, the user must also have Search access to the Parent; for a file, the user must also have Read access to the Parent Directory.	

PACKET FORMAT:



FPEnumerate

This call is used to enumerate the contents of a directory. The reply is composed of a number of file and/or directory parameter structures.

INPUTS:	SRefNum (INT)	session refnum
	VolumeID (INT)	volume identifier
	DirID (LONG)	ancestor directory identifier
	FileBitmap (INT)	bitmap describing which parameters are to be returned if the enumerated object is a file (the corresponding bit should be set). This field is the same as that in the FPGetFileDirParms call, and may be null.
	DirBitmap (INT)	bitmap describing which parameters are to be returned if the enumerated object is a directory (the corresponding bit should be set). This field is the same as that in the FPGetFileDirParms call, and may be null.
	ReqCount (INT)	maximum number of structures to be returned
	StartIndex (INT)	directory offspring index, described below
	MaxReplySize (INT)	maximum size of reply buffer
	PathType (BYTE)	indicates whether Pathname is composed of long names or short names: 1 = all pathname elements are short names 2 = all pathname elements are long names
	Pathname (STR)	pathname to desired directory
OUTPUTS:	FPErr (LONG)	
	FileBitmap (INT)	copy of input parameter
	DirBitmap (INT)	copy of input parameter
	ActCount (INT)	actual number of structures returned (zero if error)
	ActCount number of file/directory structures of the form:	
	StructLength (BYTE)	unsigned length of this structure, including these two 'header' bytes, and rounded up to the nearest even number
	File/DirFlag (BIT)	flag indicates whether structure describes a file or directory: 0 = file; 1 = directory. All other bits must be zero.
	File or Directory Parameters, packed in Bitmap order, with a trailing null BYTE if necessary to make the length of the entire structure even	
ERRORS:	ParamErr	unknown session refnum, volume identifier, directory identifier, or pathname type; bad pathname;
	DirNotFound	MaxReplySize is too small to hold even a single entry
	BitmapErr	input parameters do not point to an existing directory an attempt was made to retrieve a parameter which cannot be retrieved with this call; both bitmaps are empty
	AccessDenied	user does not have the rights listed below
	ObjectNotFound	no more offspring to enumerate
	ObjectTypeErr	input parameters pointed to a file

ALGORITHM: The server does an enumeration of the directory as specified with the input parameters: if the FileBitmap is empty, only directory offspring will be enumerated, and the StartIndex may range from 1 to the total number of directory offspring. Likewise, if the DirBitmap is empty, only file offspring will be enumerated, and the StartIndex may range from 1 to the total number of file offspring. If both Bitmaps are non-empty, the StartIndex may range from 1 to the total number of offspring, and structures for both files and directories will be returned. These structures are not returned in any particular order.

This call will complete when ReqCount structures have been inserted into the Reply packet (no partial structures will be returned), or when the Reply packet is full, or when there are no more offspring to enumerate.

The server retrieves the specified parameters for each enumerated offspring and packs them, in Bitmap order, in structures in the reply packet along with copies of the input Bitmaps inserted before the structures. In order to keep all variable-length parameters at the end of each structure (even if more parameters are later added), all such parameters like the Long Name and Short Name fields will be represented in the Bitmap order as fixed-length offsets (INTs) from the start of the parameters in each structure (not the start of the bitmap and not the start of the 'header' bytes) to the start of the variable-length fields. Each structure may be null-padded to make its length even.

A BitmapErr will be returned if an attempt is made to retrieve the DirectoryID parameter for a directory on a Variable-DirID volume.

If NoErr is returned, then all the structures returned in the Reply packet will be valid. If any error occurs, there will be no valid structures in the Reply packet.

RIGHTS:

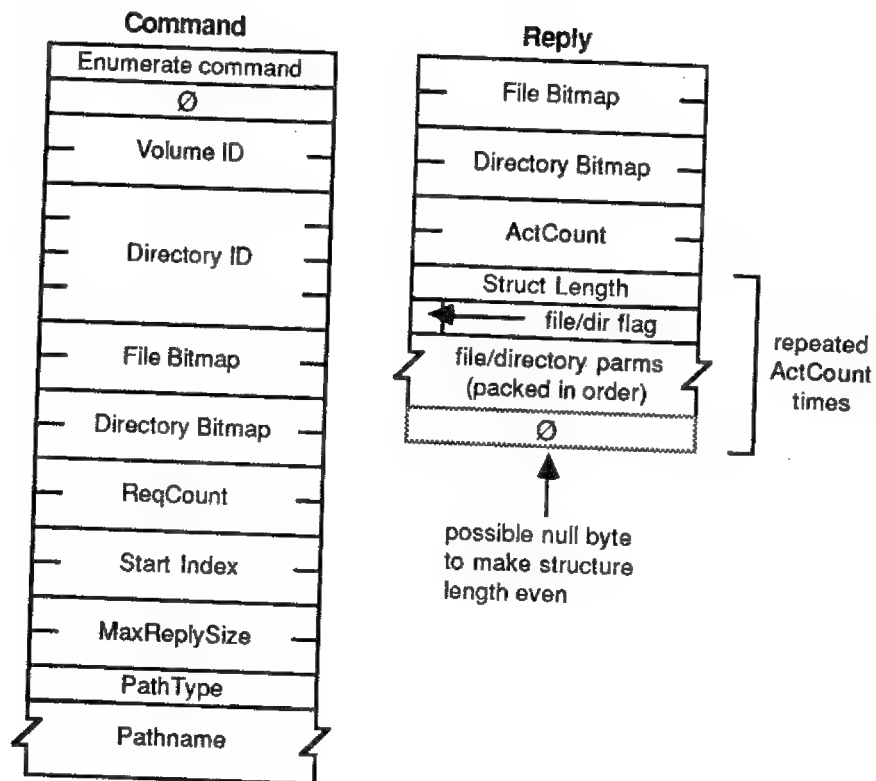
The user must have previously called FPOpenVol for this volume. In addition, the user must have Search access rights to all ancestors except this directory; Search access right to this directory is needed to be able to enumerate directory offspring. Read access right is needed to this Directory in order to be able to enumerate file offspring.

NOTES:

Since enumerating a large directory may take several calls, it is possible (since other users may be adding to or deleting from the directory) for the enumeration to miss offspring or return duplicate offspring. To be safe, keep enumerating until an ObjectNotFound error is returned, and filter out duplicate entries.

A given offspring is not guaranteed to occupy the same index number in the Parent Directory from one enumeration to the next.

PACKET FORMAT:



FPFlush

This call is used to flush to disk any data relating to the specified volume that has been modified by the user.

INPUTS: SRefNum (INT) session refnum
VolumeID (INT) volume identifier

OUTPUTS: FPErr (LONG)

ERRORS: ParamErr unknown session refnum or volume identifier

ALGORITHM: The server attempts to flush to disk as much changed information as possible. This includes (a) flushing all forks opened by the user, (b) flushing catalog information changed by the user, and (c) flushing any updated volume-level information. Since it may be difficult or impossible for all servers to guarantee that this can all be done, the above list is meant as a suggestion. Users should not rely on any or all of the above actions to actually be done.

The volume's Mod Date may change as a result of this call, but users should not rely on it since updating of the date is implementation-dependent. If no volume information was changed since the last FPFlush call, the date may or may not change.

RIGHTS: The user must have previously called FPOpenVol for this volume.

PACKET FORMAT:

Command	
Flush command	
Ø	
—	Volume ID —

FPFlushFork

Any FPWrites made to a particular file fork may be buffered by the server in order to optimize disk accesses. Within the constraints of performance, the server will try to flush (commit to disk) each file as soon as possible, yet clients can force the server to write to the disk any data buffered from previous FPWrites by issuing this call.

INPUTS: SRefNum (INT) session refnum
 OForkRefnum (INT) open fork refnum

OUTPUTS: FPErr (LONG)

ERRORS: ParamErr unknown session refnum or open fork refnum

ALGORITHM: The server will commit to disk all cached writes to the fork, and set the file's Mod Date to the server's clock if the fork was written to.

RIGHTS: No special access rights are needed to make this call.

PACKET FORMAT:

Command	
FlushFork command	
Ø	
—	OForkRefnum

FPGetAPPL

This call is used to retrieve information about a particular application from the Desktop Database.

INPUTS::

SRefNum (INT)	session refnum
DTRcfnum (INT)	Desktop Database refnum
FileCreator (RESTYPE)	creator type of the application to be returned
APPL Index (INT)	index of the APPL entry to be retrieved
Bitmap (INT)	bitmap indicating the parameters of the application file to be returned. This field is the same as the FileBitmap in the FPGetFileDirParms call.

OUTPUTS:

FPErrror (LONG)	
APPLTag (LONG)	tag information associated with the APPL entry

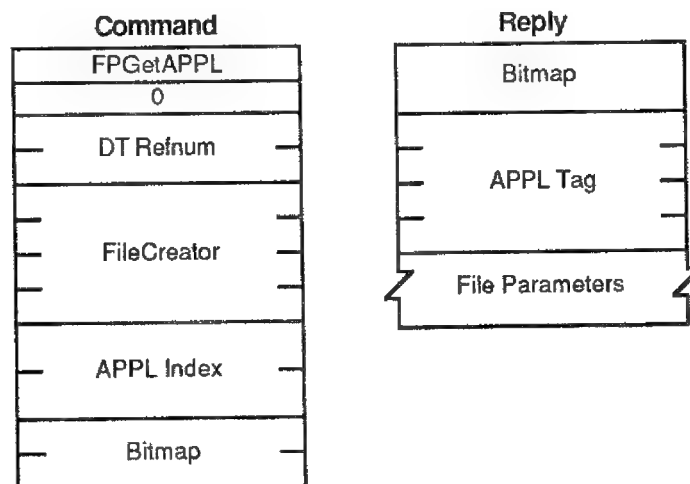
ERRORS:

ParamErr	unknown session refnum or Desktop Database refnum
ObjectNotFound	no files in the Desktop Database match input parameters
AccessDenied	user does not have the rights listed below

ALGORITHM: The entries under the specified FileCreator are examined, and the n^{th} entry, as indicated by the APPL index, is returned. Entries for applications which are not accessible by the user are not returned.

RIGHTS: The user must have previously called FPOpenDT for the corresponding volume, and must have Search access to all ancestors except the Parent and Read access to the Parent of the application whose information will be returned.

PACKET FORMAT:



FPGetComment

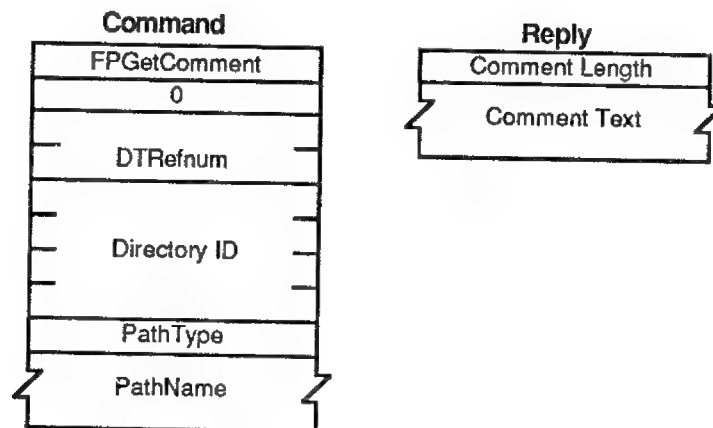
This call is used to retrieve a comment associated with a specified file or directory from the Desktop Database.

INPUTS:	SRefNum (INT)	session refnum
	DTRefnum (INT)	Desktop Database Refnum
	DirID (LONG)	directory identifier
	PathType (BYTE)	indicates whether Pathname is composed of long names or short names: 1 = all pathname elements are short names 2 = all pathname elements are long names
	Pathname (STR)	pathname to desired object
OUTPUTS:	FPErr (LONG)	
	CmtLength (BYTE)	length of the comment data
	CmtText (BYTES)	comment text
ERRORS:	ParamErr	unknown session refnum or Desktop Database refnum
	ObjectNotFound	input parameters do not point to an existing file or dir
	AccessDenied	user does not have the rights listed below
	ItemNotFound	no comment was found in the Desktop Database

ALGORITHM: The comment for the specified file is located in the Desktop Database and returned to the caller.

RIGHTS: The user must previously have called FPOpenDT for the corresponding volume. In addition, the file or directory must be present before this call is issued. If the comment is associated with a directory, the user must have Search access to all ancestors including the Parent Directory. If the comment is associated with a file, the user must have Search access to all ancestors except the Parent Directory, and Read access to the Parent Directory.

PACKET FORMAT:



FPGetFileDirParms

This call is used to retrieve parameters for an object that may be a file or a directory.

INPUTS:	SRefNum (INT)	session refnum
	VolumeID (INT)	volume identifier
	DirID (LONG)	ancestor directory identifier
	FileBitmap (INT)	bitmap describing which parameters are to be returned if the object is a file (the corresponding bit should be set):
		0 (LSB) Attributes (INT), consisting of the following flags:
		0 Invisible
		1 Multi-User
		3 DAlreadyOpen
		4 RAlreadyOpen
		5 ReadOnly
		15 Set/Clear (used in FPSetFileDirParms)
		1 Parent Directory ID (LONG)
		2 Create Date (LONG)
		3 Mod Date (LONG)
		4 Backup Date (LONG)
		5 Finder Info (32 BYTES)
		6 Long Name (INT)
		7 Short Name (INT)
		8 File Number (LONG)
		9 Data Fork Length (LONG)
		10 Resource Fork Length (LONG)
	DirBimap (INT)	bitmap describing which parameters are to be returned if the object is a directory (the corresponding bit should be set):
		0 (LSB) Attributes (INT), consisting of the following flag:
		0 Invisible
		1 Parent Directory ID (LONG)
		2 Create Date (LONG)
		3 Mod Date (LONG)
		4 Backup Date (LONG)
		5 Finder Info (32 BYTES)
		6 Long Name (INT)
		7 Short Name (INT)
		8 Directory ID (LONG)
		9 Number of Offspring (INT)
		10 Owner ID (LONG)
		11 Group ID (LONG)
		12 Access Rights (LONG), composed of the access privileges for Owner, Group, and World, and a User Rights Summary BYTE
	PathType (BYTE)	indicates whether Pathname is composed of long names or short names:
		1 = all pathname elements are short names
		2 = all pathname elements are long names
	Pathname (STR)	pathname to desired object
OUTPUTS:	FPErrror (LONG)	copy of input parameter
	FileBimap (INT)	copy of input parameter
	DirBimap (INT)	one-bit flag that indicates whether object is a file or a directory: 0 = file; 1 = directory. All other bits should be zero.
	ObjectFlag (BIT)	

Parameters

ERRORS:	ParamErr	unknown session refnum, volume identifier, or pathname type; bad pathname
	ObjectNotFound	input parameters do not point to an existing file or dir
	BitmapErr	an attempt was made to retrieve a parameter which cannot be obtained with this call
	AccessDenied	user does not have the rights listed below

ALGORITHM: The server retrieves the specified parameters for the object and packs them, in the order specified by the appropriate Bitmap, in the reply packet along with a flag indicating the type of object and a copy of the Bitmaps inserted before the parameters. In order to keep all variable-length parameters at the end of the packet (even if more parameters are later added), all such parameters like the Long Name and Short Name fields will be represented in the Bitmap order as fixed-length offsets (INTs) from the start of the parameters (not the start of the bitmap) to the variable-length fields. The actual variable-length fields are then packed after all fixed-length fields.


If the object exists but both bitmaps are null, no error will be returned. The File Bitmap, Dir Bitmap, and File/Dir Flag will be returned with no other parameters.

If the Access Rights for a directory are requested, the server will return a LONG containing the Read, Write, and Search access privileges corresponding to Owner, Group, and World. In addition, the upper byte of the Access Rights LONG is the User Rights Summary byte, indicating what privileges the user has to this directory, and whether or not the user is the owner of the directory. This Owner Bit is also set if the directory is owned by <any user>.

RIGHTS: The user must have previously called FPOpenVol for this volume. In addition, the user must have Search access rights to all ancestors except this object's Parent Directory. If the object is a directory, the user also needs Search access to the Parent Directory; else if the object is a file, the user needs Read access to the Parent Directory.

NOTES: Most Attributes are actually stored in corresponding flags within the FinderInfo field.

PACKET FORMAT:



FPGetForkParms

This call is used to retrieve parameters for a file associated with a particular open fork.

INPUTS:

SRefNum (INT)	session refnum
OForkRefnum (INT)	open fork refnum
Bitmap (INT)	bitmap describing which parameters are to be retrieved (the corresponding bit should be set). This field is the same as that in the FPGetFileDirParms call.

OUTPUTS:

FPErr (LONG)	
Bitmap (INT)	copy of the input parameter
File Parameters	

ERRORS:


ParamErr	unknown session refnum or open fork refnum
BitmapErr	an attempt was made to retrieve a parameter which cannot be obtained with this call; null bitmap
AccessDenied	fork was not opened for Read

ALGORITHM: The server retrieves the specified parameters for the file and packs them, in Bitmap order, in the reply packet. In order to keep all variable-length parameters at the end of the packet (even if more parameters are later added), all such parameters will be represented in the Bitmap order as fixed-length offsets (INTs) from the start of the parameters to the start of the variable-length fields. The actual variable-length fields are then packed after all fixed-length fields.

In AFP Version 1.1, the length of the fork indicated by the OForkRefnum may be retrieved, but a BitmapErr will be returned if an attempt is made to retrieve the length of the other fork comprising the file.

RIGHTS: The fork must have been opened for Read.

PACKET FORMAT:



FPGetIcon

This call is used to retrieve an icon from the Desktop database from a FileCreator/FileType specification.

INPUTS:

SRefNum (INT)	session refnum
DTRefnum (INT)	Desktop Database refnum
FileCreator (RESTYPE)	File's Creator type
FileType (RESTYPE)	File's type
IconType (BYTE)	Preferred icon type
Length (INT)	the number of bytes reserved for icon bitmap

OUTPUTS:

FPErrror (LONG)	
Icon Bitmap (BYTES)	The actual bitmap for the icon

ERRORS:

ParamErr	unknown session refnum or Desktop Database refnum
ItemNotFound	no icon corresponding to the input specification was found in the Desktop Database

ALGORITHM: The bitmap for the specified icon is looked up in the Desktop Database given its FileCreator, FileType, and IconType and returned to the caller if found. If no matching icon is found, an ItemNotFound error will be returned.

Note that an input length argument of zero is acceptable to test for the presence or absence of a particular icon. The size of the bitmap returned is the minimum of the requested length and the actual size of the icon.


RIGHTS: The user must have previously called FPOpenDT for the corresponding volume.

PACKET FORMAT:

Command	
FPGetIcon	
0	
DTRefNum	
FileCreator	
FileType	
IconType	
0	
Length	

FPGetIconInfo

INPUTS:	SRefNum (INT)	session refnum
	DTRefnum (INT)	Desktop Database refnum
	FileCreator (RESTYPE)	File's Creator type
	IconIndex (INT)	Index of requested icon
OUTPUTS:	FPErr (LONG)	
	IconTag (LONG)	Tag information associated with the requested icon
	FileType (RESTYPE)	The file type of the requested icon
	IconType (BYTE)	The type of the requested icon
	Size (INT)	The size of the icon bitmap
ERRORS:	ParamErr	unknown session refnum or Desktop Database refnum
	ItemNotFound	no icon corresponding to the input specification was found in the Desktop Database
ALGORITHM:	The Icon Index argument is used to determine the n th icon for the given Creator type to be returned. If the icon index is greater than the number of icons in the Desktop Database for the specified Creator type, ItemNotFound is returned.	
RIGHTS:	The user must have previously called FPOpenDT for the corresponding volume.	
PACKET FORMAT:		

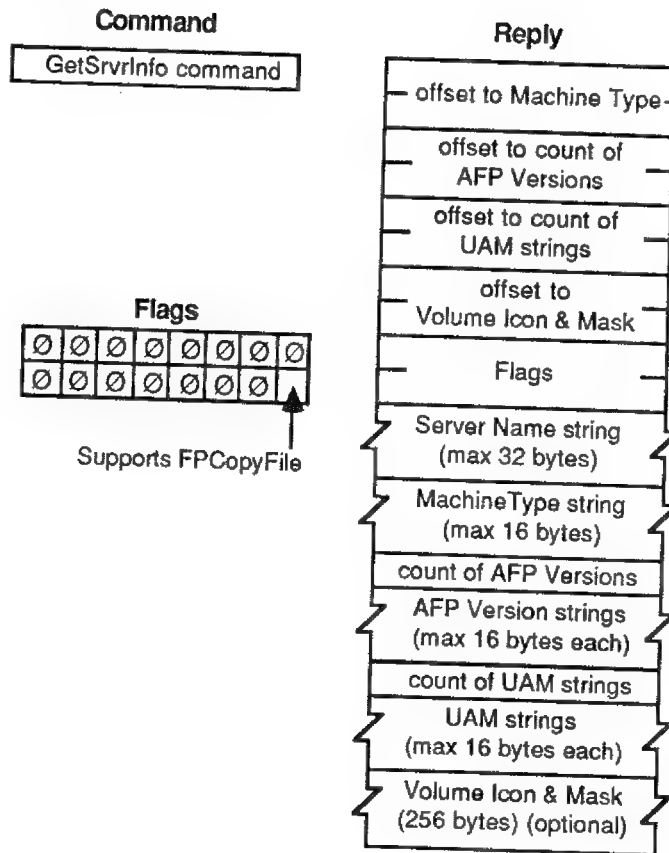


FPGetSrvrInfo

This call is used to obtain a block of descriptive information from the server, without requiring a session to be opened.

INPUTS:	SAddr (EntityAddr)	network-dependent internet address of the file server
OUTPUTS:	FPErr (LONG)	
	Flags (INT)	Flags, consisting of: Bit 0 SupportsCopyFile set if server supports the FPCopyFile call
	Server Name (STR)	the name of the server
	Machine Type (STR)	string describing the server's hardware and/or OS
	AFP Versions (STRs)	versions of AFP that the server can speak
	UAM strings (STRs)	User Authentication Methods supported by the server
	Volume Icon and Mask (256 BYTES)	
ERRORS:	NoServer	server not responding
ALGORITHM:	The info block is returned from the server. The AFP Versions and UAM strings are formatted as a one-BYTE count followed by that number of strings packed back-to-back without padding. To facilitate access to all the fields of the reply packet, the packet is formatted as shown below: the packet data begins with INT offsets to the Machine Type, AFP Versions, UAM strings, and Volume Icon and Mask. These offsets are measured relative to the start of the reply packet data. The Volume Icon and Mask field is optional; if it is not included, the Offset to Volume Icon and Mask will be zero.	
RIGHTS:	No special access rights are needed to make this call.	
NOTES:	The server may pack fields in the Reply block in any order; each field should be accessed only via the offsets (make no assumptions about how the fields are packed relative to one another). The exception to this is that the Server Name string (for which there is no offset) begins immediately after the Flags field.	

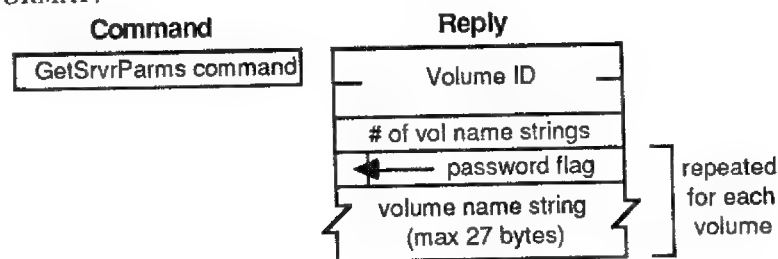
PACKET FORMAT:



FPGetSrvrParms

This call is used to retrieve server-level parameters.

INPUTS:	SRefNum (INT)	session refnum
OUTPUTS:	FPErr (LONG)	
	ServerTime (LONG)	current date-time on this server's clock
	NumVols (BYTE)	number of volumes managed by the server
	HasPassword (BIT)	flag indicating whether or not this volume is password-protected: 0 = not protected; 1 = has password (all other bits must be zero)
	VolNames (STRs)	character string names of each volume (maximum 27 bytes)
ERRORS:	ParamErr	unknown session refnum
ALGORITHM:	The volume name strings and password flags are packed together without padding in the reply.	
RIGHTS:	No special access rights are needed to make this call.	
NOTES:	This call should be implemented on a server using the ASP GetStatus mechanism.	
PACKET FORMAT:		



FPGetVolParms

This call is used to retrieve parameters for a particular volume. The volume is specified by its Volume ID as returned from the FPOpenVol call.

INPUTS:

SRefNum (INT)	session refnum
VolumeID (INT)	volume identifier
Bitmap (INT)	bitmap describing which parameters are to be returned (the corresponding bit should be set) (cannot be null):
	0 (LSB) Attributes (INT), consisting of the following flag:
	0 ReadOnly
	1 Signature (INT)
	2 Create Date (LONG)
	3 Mod Date (LONG)
	4 Backup Date (LONG)
	5 Volume ID (INT)
	6 Bytes Free (LONG) unsigned
	7 Bytes Total (LONG) unsigned
	8 Volume Name (INT)

OUTPUTS:

FPErr (LONG)	
Bitmap (INT)	copy of input parameter
Volume Parameters	

ERRORS:

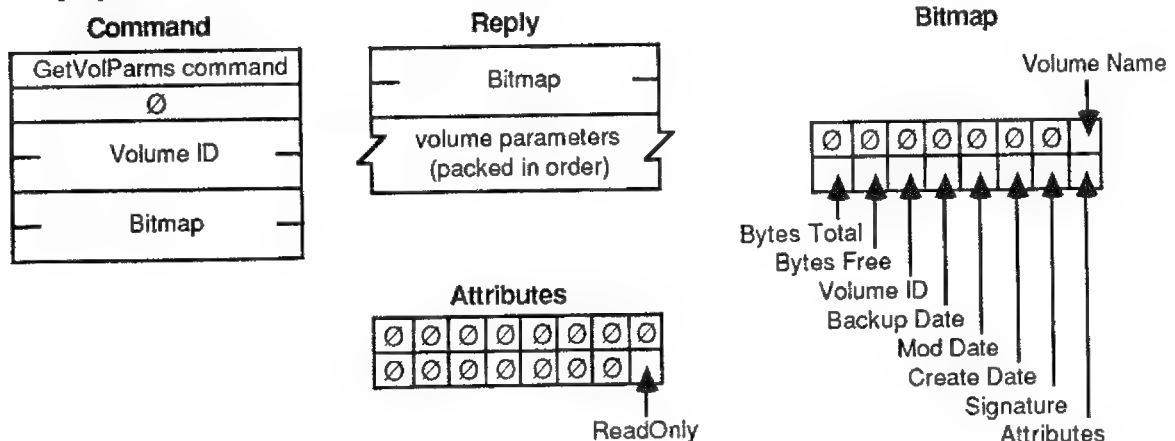
ParamErr	unknown session refnum or volume identifier
BitmapErr	an attempt was made to retrieve a parameter which cannot be obtained with this call; null bitmap

ALGORITHM: The server retrieves the specified parameters for the volume and packs them, in Bitmap order, in the reply packet along with a copy of the Bitmap inserted before the parameters. In order to keep all variable-length parameters at the end of the packet (even if more parameters are later added), all such parameters like the Volume Name field will be represented in the Bitmap order by fixed-length offsets (INTs) from the start of the parameters (not the start of the bitmap) to the start of the variable-length fields. The actual variable-length fields are then packed after all fixed-length fields.

RIGHTS: The user must have previously called FPOpenVol for this volume.

NOTES: The ReadOnly attribute is intended to be set via some administrative function, not through this protocol.

PACKET FORMAT:



FPLogin

This call is used to establish a session with a server. A protocol version is agreed upon and the user is authenticated.

INPUTS:	SAddr (EntityAddr)	network-dependent internet address of the file server
	AFPVersion (STR)	a string indicating which AFP Version is to be used
	UAM (STR)	a string indicating which User Authentication Method is to be used to authenticate the user
	UserAuthInfo (BUF)	information required to authenticate the user, dependent on the method used (may be null)

OUTPUTS:	FPErr (LONG)	session refnum to be used to refer to this session in all subsequent calls (valid if NoErr or AuthContinue error returned)
	SRefNum (INT)	
	IDNumber (INT)	returned in certain UserAuthenticationMethods to be presented in the next FPLoginCont call (only valid if AuthContinue error returned)
	UserAuthInfo (BUF)	returned in certain UserAuthenticationMethods (only valid if AuthContinue error returned)

ERRORS:	NoServer	server not responding
	BadVersNum	server cannot speak the specified AFP version
	BadUAM	unknown UserAuthenticationMethod
	ParamErr	unknown User
	UserNotAuth	UserAuthenticationMethod failed
	AuthContinue	authentication not yet complete
	ServerGoingDown	the server is in the process of shutting down
	MiscErr	user is already authenticated

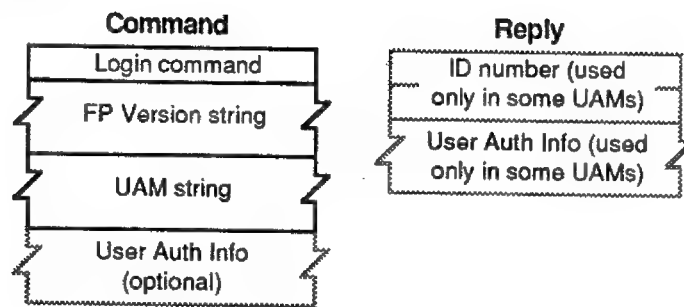
ALGORITHM: The AFP Version string, indicating which AFP Version to use, and the User Authentication Method string, indicating which UAM is to be used in authenticating the user, are sent to the server. These are packed into the command packet with no padding. Depending on which method is used, the command packet sent to the server may contain additional information like user name and password. This extra information is passed to AFP as UserAuthInfo. If the server knows how to execute that UserAuthenticationMethod, it will do so and return a UserNotAuthenticated error if that method fails. Depending on which UAM is used, there may or may not be a null byte padded between the UAM and UserAuthInfo fields. See Appendix A for more details.

Note that some UserAuthenticationMethods may return some Reply data to this call. "No User Authent" and "Cleartxt passwd" do not. Some methods like "Randnum exchange" will return UserAuthInfo data and will require an additional exchange of packets as well. In such cases, an AuthContinue error will be returned to this call, indicating that subsequent FPLoginCont calls are needed. (See Appendix A for more details.)

NOTES: If any error (other than AuthContinue) is returned, the session will not be opened.

RIGHTS: No special access rights are needed to make this call.

PACKET FORMAT:



FPLoginCont

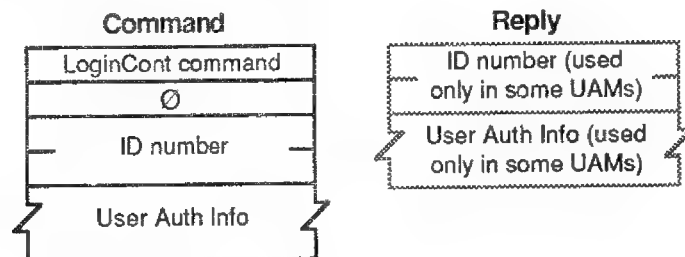
This call is used to continue the Login and authentication process with a server.

INPUTS:	SRefNum (INT)	session refnum
	IDNumber (INT)	number returned from the previous FPLogin or FPLoginCont call
	UserAuthInfo (BUF)	information required to authenticate the user, dependent on the method used
OUTPUTS:	FPErr (LONG)	returned in certain UserAuthenticationMethods to be presented in the next FPLoginCont call (only valid if AuthContinue error returned)
	IDNumber (INT)	returned in certain UserAuthenticationMethods (only valid if AuthContinue error returned)
ERRORS:	NoServer	server not responding
	UserNotAuth	UserAuthenticationMethod failed
	AuthContinue	authorization not yet complete

ALGORITHM: The ID number and UserAuthInfo are sent to the server, which uses them to execute the next step in the authentication method. If an additional exchange of packets is required, an AuthContinue error will be returned. Otherwise, either NoErr (meaning the user has been authenticated) or UserNotAuth (meaning the authentication method has failed) will be returned. If NoErr, a valid SRefNum will be returned for use in subsequent calls. If UserNotAuth, the session is closed by the server and the SRefnum is invalidated.

RIGHTS: No special access rights are needed to make this call.

PACKET FORMAT:



FPLogout

This call is used to terminate a session with a server

INPUTS: SRefNum (INT) session refnum

OUTPUTS: FPErr (LONG)

ERRORS: ParamErr unknown session refnum

ALGORITHM: The server flushes and closes any forks opened by this session, frees up all session-related resources and invalidates the session refnum.

RIGHTS: No special access rights are needed to make this call.

PACKET FORMAT:


Command

Logout command

FPMaPID

This call is used to map a User ID to a User Name, or a Group ID to a Group Name.

INPUTS:	SRefNum (INT)	session refnum
	Function (BYTE)	function code: 1 = map User ID to User Name 2 = map Group ID to Group Name
	ID (LONG)	item to be mapped, either User ID or Group ID
OUTPUTS:	FPErr (LONG)	
	Name (STR)	name corresponding to input ID
ERRORS:	ParamErr	unknown session refnum or function code; no ID was passed in command packet
	ItemNotFound	ID not recognized
ALGORITHM:	The server attempts to find the Creator Name or Group Name corresponding to the specified Creator ID or Group ID. An ItemNotFound error is returned if the ID does not exist in the server's list of valid User or Group IDs.	
RIGHTS:	No special access rights are needed to make this call.	
NOTES:	A User ID or Group ID of zero will map to the null string.	
PACKET FORMAT:		



FPMapName

This call is used to map a User Name to a User ID, or a Group Name to a Group ID.

INPUTS:

SRefNum (INT)	session refnum
Function (BYTE)	function code: 3 = map User Name to User ID 4 = map Group Name to Group ID
Name (STR)	item to be mapped, either User Name or Group Name

OUTPUTS:

FPErr (LONG)	
ID (LONG)	ID corresponding to input Name

ERRORS:


ParamErr	unknown session refnum or function code
ItemNotFound	name not recognized

ALGORITHM: The server attempts to find the User ID or Group ID corresponding to the specified User Name or Group Name. An ItemNotFound error is returned if the name does not exist in the server's list of valid User or Group names.

RIGHTS: No special access rights are needed to make this call.

NOTES: A null User or Group Name will map to an ID of zero.

PACKET FORMAT:

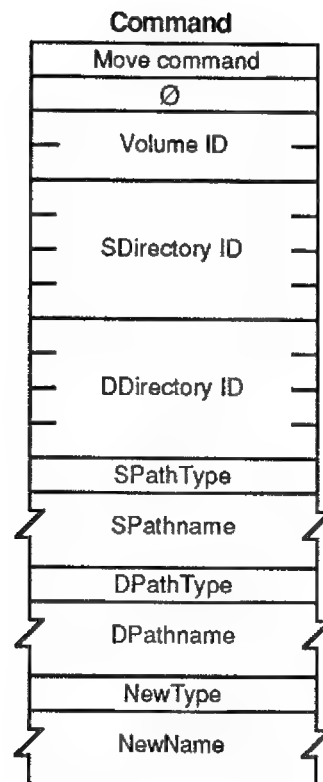


FPMove

This call is used to move (not just copy) a directory or file to another location on a single volume (source and destination must be on the same volume). An object cannot be moved from one volume to another with this call, even though both volumes may be managed by the server. The destination of the move is specified by providing a DirID and Pathname that indicate the object's new Parent Directory.

INPUTS:	SRefNum (INT)	session refnum
	VolumeID (INT)	volume identifier
	SDirID (LONG)	source ancestor directory identifier
	SPathType (BYTE)	indicates whether SPathname is composed of long names or short names: 1 = all pathname elements are short names 2 = all pathname elements are long names
	SPathname (STR)	pathname of file or directory to be moved (may be null if a directory is to be moved)
	DDirID (LONG)	destination ancestor directory identifier
	DPathType (BYTE)	indicates whether DPathname is composed of long names or short names (same values as SPathType)
	DPathname (STR)	pathname to the destination Parent Directory (may be null)
	NewType (BYTE)	indicates whether NewName is a long name or a short name (same values as SPathType)
	NewName (STR)	new name of file or directory (may be null)
OUTPUTS:	FPErr (LONG)	
ERRORS:	ParamErr	unknown session refnum, volume identifier, or pathname type; bad pathname or NewName
	ObjectNotFound	input parameters do not point to an existing file or directory
	ObjectExists	a file or directory with the name NewName already exists
	CantMove	an attempt was made to move a directory into one of its descendent directories
	AccessDenied	user does not have the right to move the file/directory
ALGORITHM:	If the object to be moved is a directory, the directory and all its descendents will be moved. The file or directory is moved (deleted from its original Parent Directory) and renamed to its new name. The creation of Long and Short names is performed as described in Appendix B. The object's Mod Date, and the Mod Date of the object's Parent Directory are set to the server's clock. If NewName is null, the object will not be renamed. Its Parent Directory ID will be set to the destination Parent DirID, but all other parameters remain unchanged. The parameters of all descendent directories and files remain unchanged.	
RIGHTS:	The user must have previously called FPOpenVol for the volume. To move a directory, the user must have Search access rights to all ancestors down to and including the source and destination Parents, as well as Write access right to those directories. To move a file, Search access rights are needed for all ancestors except the source and destination Parents, as well as Read and Write access rights to the source Parent and Write access right to the destination Parent.	

PACKET FORMAT:

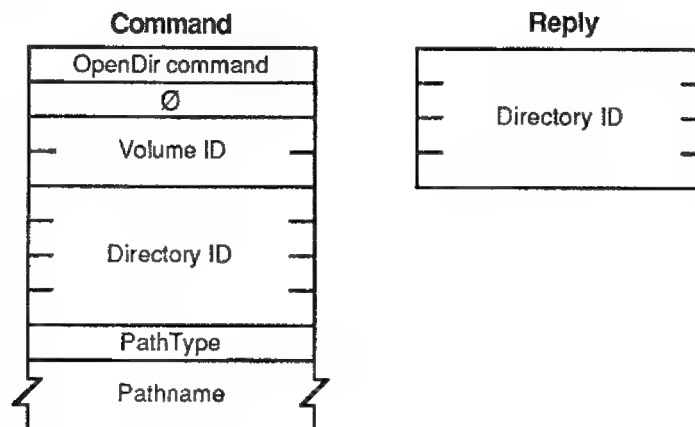


FPOpenDir

This call is used to open a directory on a Variable-DirID volume and obtain its directory identifier.

INPUTS:	SRefNum (INT)	session refnum
	VolumeID (INT)	volume identifier
	DirID (LONG)	ancestor directory identifier
	PathType (BYTE)	indicates whether Pathname is composed of long names or short names: 1 = all pathname elements are short names 2 = all pathname elements are long names
	Pathname (STR)	pathname to desired directory (cannot be null)
OUTPUTS:	FPErrror (LONG)	
	DirID (LONG)	identifier of specified directory
ERRORS:	ParamErr	unknown session refnum, volume identifier, or pathname type; bad pathname
	ObjectNotFound	input parameters do not point to an existing directory
	AccessDenied	user does not have the rights listed below
	ObjectTypeErr	input parameters point to a file
ALGORITHM:	If Volume ID specifies a Variable-DirID volume, the server will generate a variable Directory ID for the directory specified by the other input parameters. If the volume is of Fixed-DirID type, the server will return the fixed DirectoryID belonging to this directory.	
RIGHTS:	The user must have previously called FPOpenVol for this volume. In addition, the user must have Search access rights to all ancestors down to and including this directory's Parent Directory.	
NOTES:	This call must be issued to obtain a DirID for a directory and to subsequently access the directory on a Variable-DirID volume. It is not considered an error to invoke this call for directories on other types of volumes, although this is not the recommended way to obtain the parameter. Use the FPGetFileDirParms or FPEnumerate calls instead.	

PACKET FORMAT:



FPOpenDT

This call is used to retrieve an icon from the Desktop database from a FileCreator/FileType specification.

INPUTS: DTRefNum (INT) Desktop Database RefNum


OUTPUTS: FPErrror (LONG)
DTRefNum (INT) Refnum for use on future Desktop Manager calls

ERRORS: ParamErr unknown session refnum or volume identifier

ALGORITHM: The desktop Database on the selected volume is opened, and a refnum (unique among all volumes on the server) is returned for use in subsequent calls.

RIGHTS: No special rights are needed to make this call.

PACKET FORMAT:



FPOpenFork

This call is used to open the data or resource fork of an existing file for the purpose of reading from it or writing to it. Each fork must be opened separately; a unique OpenForkRefnum will be returned for each.

INPUTS:	SRefNum (INT)	session refnum
	VolumeID (INT)	volume identifier
	DirID (LONG)	ancestor directory identifier
	Bitmap (INT)	bitmap describing which parameters are to be returned (the corresponding bit should be set). This field is the same as that in the FPGetFileDirParms call. (can be null)
	AccessMode (INT)	desired access and sharing modes, as specified by any combination of the following bits: 0 Read - allows the file to be read 1 Write - allows the file to be written to 4 DenyRead - denies others the right to read the fork while this user has it open 5 DenyWrite - denies others the right to write to the fork while this user has it open See Appendix C for a detailed explanation of the use of the Deny bits.
	PathType (BYTE)	indicates whether Pathname is composed of long names or short names 1 = all pathname elements are short names 2 = all pathname elements are long names
	Pathname (STR)	pathname to desired file; cannot be null
	Rsrc/DataFlag (BIT)	flag to indicate which fork is to be opened (all other bits must be zero): 0 = data fork 1 = resource fork
OUTPUTS:	FPErr (LONG)	copy of input parameter
	Bitmap (INT)	
	OForkRefnum	refnum used to refer to this fork in subsequent calls
	File Parameters	
ERRORS:	ParamErr	unknown session refnum, volume identifier, or pathname type; null or bad pathname
	ObjectNotFound	input parameters do not point to an existing file
	BitmapErr	an attempt was made to retrieve a parameter which cannot be obtained with this call (file will not be opened)
	DenyConflict	fork cannot be opened because Deny modes conflict (parameters will be returned)
	AccessDenied	user does not have the rights listed below
	ObjectTypeErr	input parameters point to a directory
	TooManyFilesOpen	the server cannot open another fork

ALGORITHM: The server opens the specified fork if the user has the proper access rights for the requested Access Mode, and if the requested Access Mode does not conflict with already-open access paths to this fork.

If the open is successful, the server retrieves the specified parameters for the file and packs them, in Bitmap order, in the reply packet along with a copy of the Bitmap and an OForkRefnum inserted before the parameters. This OForkRefnum is to be used in all subsequent calls involving the open fork.

File parameters will be returned only if the call completes with no error or with a Deny Conflict error. In the latter case, the server will return zero for the OForkRefnum but valid parameters as requested so that the user can determine if he is the one who already has the fork open.

In order to keep all variable-length parameters at the end of the packet (even if more parameters are later added), all such parameters like the Long Name and Short Name fields must be represented in the Bitmap order as fixed-length offsets (INTs) from the start of the parameters (not the start of the bitmap) to the variable-length fields. The actual variable-length fields are then packed after all fixed-length fields.

RIGHTS:

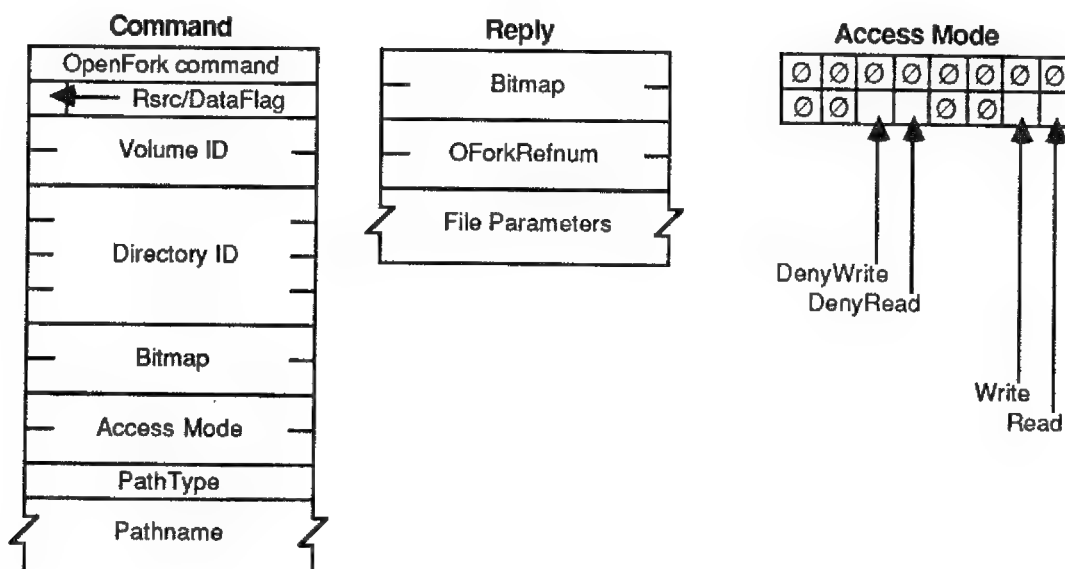
The user must have previously called FPOpenVol for this volume. To open a fork for Read, or in the special case in which neither Read nor Write access is requested, the user must have Search access to all ancestors except the Parent, as well as Read access to the Parent.

To open the fork for Write, the volume must not be marked Read Only. If both forks are currently empty, the user must have Search or Write access to all ancestors except the Parent, as well as Write access to the Parent. If either fork is non-empty and one of them is being opened for Write, the user must have Search access to all ancestors except the Parent, as well as Read and Write access to the Parent.

NOTES:

If the fork was successfully opened and the user requested the file's Attributes in the Bitmap, the appropriate DAlreadyOpen or RAlreadyOpen bits will be set.

PACKET FORMAT:

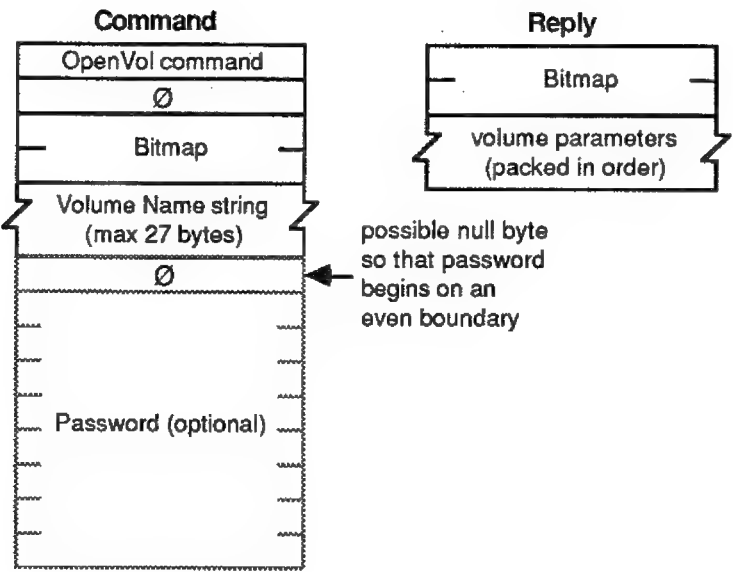


FPOpenVol

This call is used to "mount" a volume. It must be called once before any other call can be made to access objects on the volume.

INPUTS:	SRefNum (INT)	session refnum
	Bitmap (INT)	bitmap describing which parameters are to be returned (the corresponding bit should be set). This field is the same as that in the FPGetVolParms call (cannot be null)
	VolumeName (STR)	name of the volume as returned by the FPGetSrvrParms call (maximum 27 bytes)
	Password (8 BYTES)	optional password
OUTPUTS:	FPErr (LONG)	
	Bitmap (INT)	copy of input parameter
	Volume Parameters	
ERRORS:	ParamErr	unknown session refnum or volume name
	BitmapErr	an attempt was made to retrieve a parameter which cannot be obtained with this call; null bitmap
	AccessDenied	Password not supplied or does not match
ALGORITHM:	The password is sent in the command packet in clear text, padded (suffixed) with null bytes to its full 8-byte length. Password comparison is case sensitive.	
	If the volume is password-protected, the server will check that the Password supplied by the user matches the one kept with the volume. If they do not match, or if no Password was supplied, an AccessDenied error will be returned.	
	If the Passwords match, or if the volume is not password-protected, the server will retrieve the requested parameters and pack them into the reply packet. The server will mark the volume as "mounted", meaning that this user has permission to make calls relating to objects on this volume.	
RIGHTS:	No special access rights are needed to make this call.	
NOTES:	This call cannot be made with a null Bitmap; the Bitmap should at least request the Volume ID to be returned as there is no other way to retrieve this parameter and it is needed for most subsequent calls.	
	FPOpenVol can be called multiple times without an intervening FPCloseVol call; however, a single FPCloseVol call will invalidate the VolID.	

PACKET FORMAT:

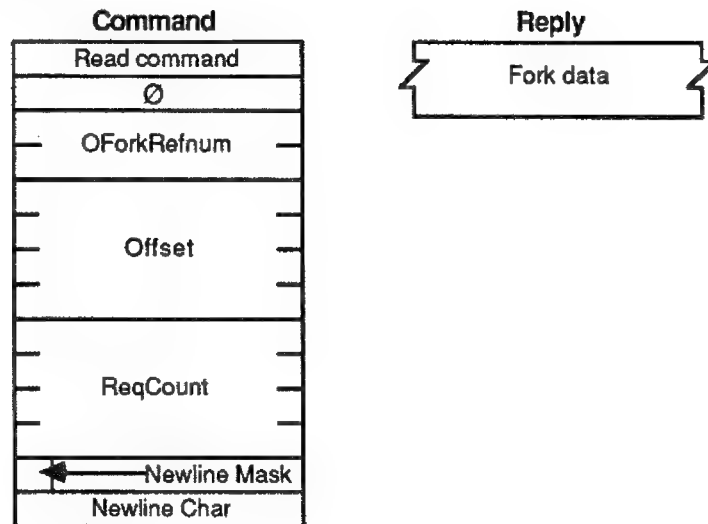


FPRead

This call is used to read a block of data from an open fork.

INPUTS:	SRefNum (INT)	session refnum
	OForkRefnum (INT)	open fork refnum
	Offset (LONG)	byte offset from beginning of fork at which the read is to begin
	ReqCount (LONG)	number of bytes to read
	NewlineMask (BYTE)	mask which is ANDed with each character and compared with the NewlineChar to test for the end of the block. In AFP Version 1.1, only two masks are valid. \$0= ignore NewlineChar \$FF= a NewlineChar is specified
	NewlineChar (BYTE)	any ASCII character from \$00 to \$FF inclusive which, when encountered in reading the fork (after masking), causes the read operation to terminate
OUTPUTS:	FPErr (LONG)	
	ActCount (LONG)	number of bytes actually read from the fork
	Fork data	
ERRORS:	ParamErr	unknown session refnum or open fork refnum; negative ReqCount or Offset; invalid Newline mask
	AccessDenied	fork was not opened for Read
	EOFErr	end-of-fork was reached
	LockErr	some or all of requested range is locked by another user
ALGORITHM:	The fork is read starting Offset bytes from the beginning of the fork and terminating at the first NewlineChar encountered (if NewlineMask is set to \$FF), the end-of-fork, the start of a locked range, or when ReqCount bytes have been read. If the end-of-fork (or the start of a locked range) was reached, all data read up to the fork end (the start of the locked range) will be returned, and an EOFErr (LockErr) will be returned.	
	Reading a byte that was never written to the fork will return an undefined value.	
RIGHTS:	The fork must have been opened for Read.	
NOTES:	Locking of the range should be done prior to this call since the underlying transport mechanism may force the request to be broken up into multiple smaller requests. Although the range may not be locked when the call begins execution, it is possible for another user to lock some or all of the range before this call completes, causing the Read to succeed partially.	

PACKET FORMAT:



FPRemoveAPPL

This call is used to remove an APPL mapping from the Desktop Database.

INPUTS:

SRefNum (INT)	session refnum
DTRefnum (INT)	Desktop Database refnum
DirID (LONG)	Directory identifier
FileCreator (RESTYPE)	Creator type of application to be removed
PathType (BYTE)	indicates whether Pathname is composed of long names or short names 1 = all pathname elements are short names 2 = all pathname elements are long names
PathName (STRING)	pathname to the application being removed

OUTPUTS: FPErr (LONG)

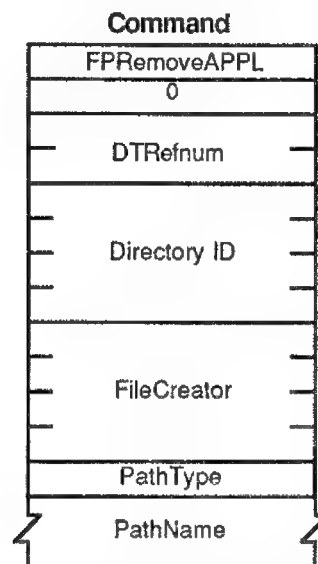
ERRORS:

ParamErr	unknown session refnum or Desktop Database refnum
ObjectNotFound	input parameters do not point to an existing file
AccessDenied	user does not have the rights listed below
ItemNotFound	no APPL entry corresponding to the input parameters was found in the Desktop Database

ALGORITHM: The entry for the application specified is located in the Desktop /Database using the FileCreator information. If an entry is found for the specified Directory ID/File name, the entry is removed.

RIGHTS: The user must have previously called FPOpenDT for the corresponding volume. In addition, the file must be present in the specified directory before this call is issued. The user must have Search access to all ancestors except the Parent, as well as Read and Write access to the Parent.

PACKET FORMAT:



FPRemoveComment

This call is used to remove a comment from the Desktop Database

INPUTS:

SRefNum (INT)	session refnum
DTRefnum (INT)	Desktop Database refnum
DirID (LONG)	directory identifier
PathType (BYTE)	indicates whether Pathname is composed of long names or short names 1 = all pathname elements are short names 2 = all pathname elements are long names
PathName (STRING)	the pathname for the file or folder associated with the comment

OUTPUTS: FPErr (LONG)


ERRORS:

ParamErr	unknown session refnum, volume identifier, or pathname type; bad pathname
ItemNotFound	no comment found in Desktop Database
AccessDenied	user does not have the rights listed below
ObjectNotFound	input parameters do not point to an existing object

ALGORITHM: The comment associated with the file or folder specified is removed from the Desktop Database.

RIGHTS: The user must have previously called FPOpen DT for the corresponding volume. If the comment is associated with a non-empty directory, the user must have Search access to all ancestors including the Parent Directory, plus Write access to the Parent. If the comment is associated with an empty directory, the user must have Search or Write access to all ancestors including the Parent Directory, plus Write access to the Parent. If the comment is associated with a non-empty file, the user must have Search access to all ancestors except the Parent Directory, plus Read and Write access to the Parent. If the comment is associated with an empty file, the user must have Search or Write access to all ancestors except the Parent Directory, plus Write access to the Parent.

PACKET FORMAT:

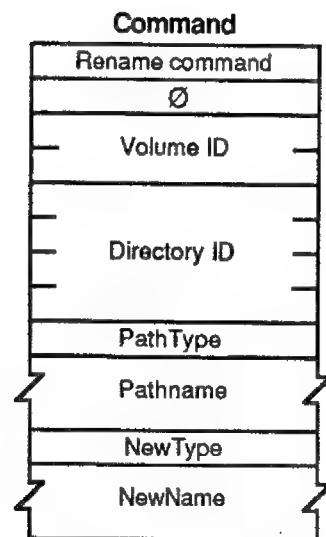


FPRename

This call is used to rename either a directory or a file.

INPUTS:	SRefNum (INT)	session refnum
	VolumeID (INT)	volume identifier
	DirID (LONG)	ancestor directory identifier
	PathType (BYTE)	indicates whether Pathname is composed of long names or short names. 1 = all pathname elements are short names 2 = all pathname elements are long names
	Pathname (STR)	pathname of file or directory to be renamed (may be null if a directory is to be renamed)
	NewType (BYTE)	indicates whether NewName is a long name or short name, (same values as PathType)
	NewName (STR)	new name of file or directory (cannot be null)
OUTPUTS:	FPErr (LONG)	
ERRORS:	ParamErr	unknown session refnum, volume identifier, or pathname type; bad pathname or NewName
	ObjectNotFound	input parameters do not point to an existing file or directory
	ObjectExists	a file or directory with the name NewName already exists
	AccessDenied	user does not have the right to rename the file/directory
	CantRename	an attempt was made to rename a volume
ALGORITHM:	The object is renamed to its new name. The creation of Long and Short names is performed as described in Appendix B. The Mod Date of the Parent Directory is set to the server's clock.	
RIGHTS:	The user must have previously called FPOpenVol for the volume. In addition, the user must have Search access rights to all ancestors except the object's Parent Directory, as well as Write access right to the Parent directory. To rename a directory, the user must also have Search access to the Parent; else to rename a file, the user must also have Read access to the Parent.	

PACKET FORMAT:



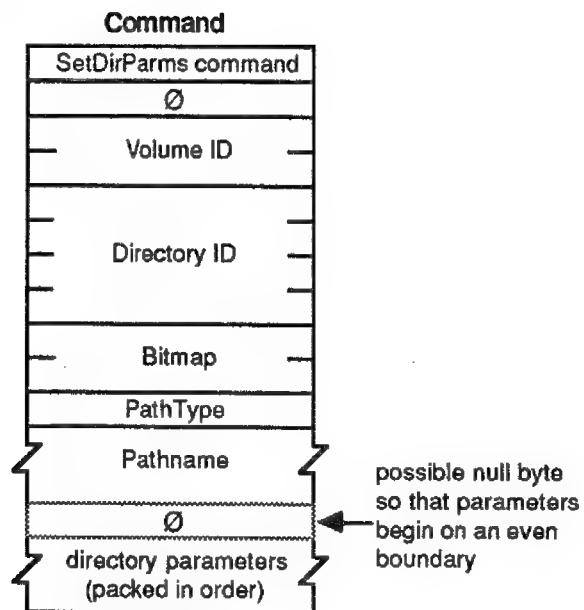
FPSetDirParms

This call is used to set parameters for a particular directory.

INPUTS:	SRefNum (INT)	session refnum
	VolumeID (INT)	volume identifier
	DirID (LONG)	ancestor directory identifier
	Bitmap (INT)	bitmap describing which parameters are to be set (the corresponding bit should be set). This field is the same as that in the FPGetFileDirParms call.
	PathType (BYTE)	indicates whether Pathname is composed of long names or short names: 1 = all pathname elements are short names 2 = all pathname elements are long names
	Pathname (STR)	pathname to desired directory
	Directory Parameters	
OUTPUTS:	FPError (LONG)	
ERRORS:	ParamErr	unknown session refnum, volume identifier, or pathname type;
		bad pathname; owner or group ID not valid
	ObjectNotFound	input parameters do not point to an existing directory
	BitmapErr	an attempt was made to set a parameter which cannot be set with this call; null bitmap
	AccessDenied	user does not have the rights listed below
	ObjectTypeErr	input parameters point to a file
ALGORITHM:	The server sets the specified parameters for the directory which must be packed, in Bitmap order, in the command packet. In order to keep all variable-length parameters at the end of the packet (even if more parameters are later added), all such parameters like the Long Name and Short Name fields must be represented in the Bitmap order as fixed-length offsets (INTs) from the start of the parameters to the start of the variable-length fields. The actual variable-length fields are then packed after all fixed-length fields.	
	Changing a directory's AccessRights will immediately affect other currently open sessions.	
	If the Access Controls or Owner ID or Group ID are set in this call, and the Mod Date is not set, the Mod Date will be set to the the server's clock.	
RIGHTS:	The user must have previously called FPOpenVol for this volume. To set a directory's Access Controls, Owner ID, or Group ID, the user must have Search or Write access rights to all ancestors including this directory's Parent Directory, and must be the owner of the directory. To set any other parameter for an empty directory, the user must have Search or Write access to all ancestors except the Parent, as well as Write access to the Parent. To set any other parameter for a non-empty directory, the user must have Search access to all ancestors including the Parent, as well as Write access to the Parent.	
	If the user lacks the access rights to set any one of a number of parameters, an Access Denied error will be returned and no parameters will be set.	
NOTES:	This call cannot be used to set a directory's name (use FPRename), Parent Directory (use FPMove), Directory ID, or number of offspring.	
	The user needs to be the owner of the directory to set the directory's Access Controls, Owner ID, or Group ID only.	

A null byte may be added between the Pathname and the Directory Parameters so that the Parameters begin on an even boundary in the Command Buffer.

PACKET FORMAT:



FPSetFileParms

This call is used to set parameters for a particular file.

INPUTS:	SRefNum (INT)	session refnum
	VolumeID (INT)	volume identifier
	DirID (LONG)	ancestor directory identifier
	Bitmap (INT)	bitmap describing which parameters are to be set (the corresponding bit should be set). This field is the same as that in the FPGetFileDirParms call.
	PathType (BYTE)	indicates whether Pathname is composed of long names or short names: 1 = all pathname elements are short names 2 = all pathname elements are long names
	Pathname (STR)	pathname to desired file
	File Parameters	
OUTPUTS:	FPErr (LONG)	
ERRORS:	ParamErr	unknown session refnum, volume identifier, or pathname type; null or bad pathname
	ObjectNotFound	input parameters do not point to an existing file
	AccessDenied	user does not have the right to access the file
	BitmapErr	an attempt was made to set a parameter which cannot be set with this call; null bitmap
	ObjectTypeErr	input parameters point to a directory

ALGORITHM: The server sets the specified parameters for the file which must be packed, in Bitmap order, in the command packet. In order to keep all variable-length parameters at the end of the packet (even if more parameters are later added), all such parameters must be represented in the Bitmap order as fixed-length offsets (INTs) from the start of the parameters to the start of the variable-length fields. The actual variable-length fields are then packed after all fixed-length fields.

In AFP Version 1.1, only the following parameters may be set or cleared: Attributes (all except DAlreadyOpen and RAlreadyOpen), Create Date, Last Mod Date, Backup Date, and FinderInfo.


If the Attributes field is included, the Set/Clear bit is used to indicate that the specified Attributes (whose corresponding bits are set) are to be either set or cleared (0 = clear specified Attributes, 1 = set specified Attributes). Hence it is not possible to set some Attributes and clear others in the same call.

RIGHTS: The user must have previously called FPOpenVol for this volume. If the file is empty (both forks are of zero length), the user must have Search or Write access rights to all ancestors except this file's Parent Directory, as well as Write access right to the Parent Directory. If either fork is non-empty, the user must have Search access to all ancestors except the Parent, as well as Read and Write access to the Parent.

NOTES: This call cannot be used to set a file's name (use FPRename), Parent Directory (use FPMove), File Number, or fork lengths.

A null byte may be added between the Pathname and the File Parameters so that the Parameters begin on an even boundary in the Command Buffer.

PACKET FORMAT:



FPSetFileDirParms

This call is used to set parameters for a particular object, either a file or a directory.

INPUTS:	SRefNum (INT)	session refnum
	VolumeID (INT)	volume identifier
	DirID (LONG)	ancestor directory identifier
	Bitmap (INT)	bitmap describing which parameters are to be set (the corresponding bit should be set). This field is the same as the FileBitmap or DirBitmap in the FPGetFileDirParms call (only the parameters which are common to both bitmaps may be set by this call).
	PathType (BYTE)	indicates whether Pathname is composed of long names or short names, or unknown: 1 = all pathname elements are short names 2 = all pathname elements are long names
	Pathname (STR) Parameters	pathname to desired object
OUTPUTS:	FPErr (LONG)	
ERRORS:	ParamErr	unknown session refnum, volume identifier, or pathname type; bad pathname
	ObjectNotFound	input parameters do not point to an existing object
	AccessDenied	user does not have the rights listed below
	BitmapErr	an attempt was made to set a parameter which cannot be set with this call; null bitmap

ALGORITHM: In AFP Version 1.1, the only parameters that may be set or cleared by this call are: Attributes (only the Invisible attribute), Create Date, Last Mod Date, Backup Date, and FinderInfo. Note that these parameters are common to both files and directories. The server sets the specified parameters for the object which must be packed, in Bitmap order, in the command packet. In order to keep all variable-length parameters at the end of the packet (even if more parameters are later added), all such parameters must be represented in the Bitmap order as fixed-length offsets (INTs) from the start of the parameters to the start of the variable-length fields. The actual variable-length fields are then packed after all fixed-length fields.

If the Attributes field is included, the Set/Clear bit is used to indicate that the specified Attributes (whose corresponding bits are set) are to be either set or cleared (0 = clear specified Attributes, 1 = set specified Attributes). Hence it is not possible to set some Attributes and clear others in the same call.

RIGHTS: The user must have previously called FPOpenVol for this volume. To set the parameters for a non-empty directory, the user needs Search access to all ancestors including the Parent Directory, as well as Write access to the Parent. To set the parameters for an empty directory, the user needs Search or Write access to all ancestors except the Parent Directory, as well as Write access to the Parent.

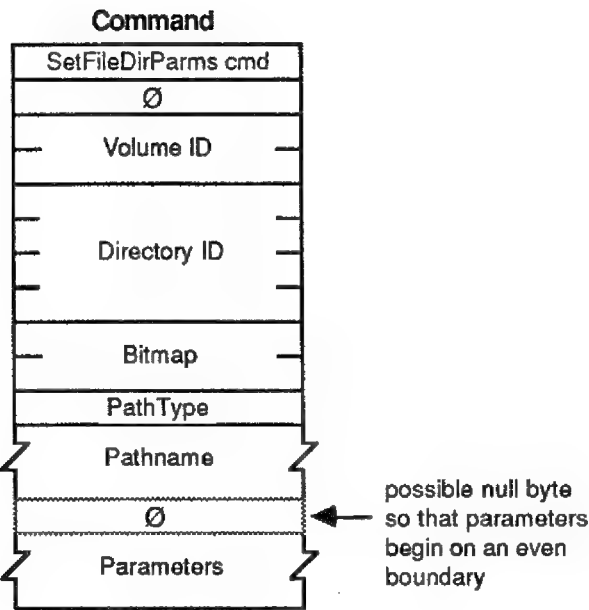
To set the parameters for a non-empty file, the user needs Search access to all ancestors except the Parent Directory, as well as Read and Write access to the Parent. To set the parameters for an empty file, the user needs Search or Write access to all ancestors except the Parent Directory, as well as Write access to the Parent.

NOTES: The Create Date, Mod Date, Backup Date, and FinderInfo may be set with the FPSetFileParms or FPSetDirParms calls if the user knows that the object is a file or a directory. To set a directory's Access Rights, Creator ID, or Group ID, use the

FPSetDirParms call. To set a file's Attributes other than Invisible, use the **FPSetFileParms** call.

A null byte may be added between the Pathname and the Parameters so that the Parameters begin on an even boundary in the Command Buffer.

PACKET FORMAT:



FPSetForkParms

This call is used to set parameters for a file associated with a particular open fork.

INPUTS:

SRefNum (INT)	session refnum
OForkRefnum (INT)	open fork refnum
Bitmap (INT)	bitmap describing which parameters are to be set (the corresponding bit should be set). This field is the same as that in the FPGetFileDirParms call; however, in AFP Version 1.1, only the fork length may be set.

Fork Length

OUTPUTS: FPErr (LONG)

ERRORS:

ParamErr	unknown session refnum or open fork refnum
BitmapErr	an attempt was made to set a parameter which cannot be set with this call; null bitmap
DiskFull	no more space on the volume
LockErr	locked range conflict
AccessDenied	fork was not opened for Write


ALGORITHM: The Bitmap and fork length are passed to the server, which changes the length of the fork specified by OForkRefnum. A BitmapErr will be returned if an attempt is made to set the length of the other fork comprising the file, or any other file parameter.

A LockErr will be returned if one of the following conditions occurs: a) the user attempts to set the end-of-fork into a range locked by another user, B) the file is truncated so as to eliminate a range locked by another user, c) the range between the old end-of-fork and the new end-of-fork intersects any lock owned by another user.

RIGHTS: The fork must have been opened for Write.

NOTES: This call cannot be used to set a file's name (use FPRename), Parent Directory (use FPMove), or File Number.

PACKET FORMAT:



FPSetVolParms

This call is used to set the parameters for a particular volume. The volume is specified by its VolumeID as returned from the FPOpenVol call.

INPUTS:

SRefNum (INT)	session refnum
VolumeID (INT)	volume identifier
Bitmap (INT)	bitmap describing which parameters are to be set (the corresponding bit should be set). This field is the same as that in the FPGetVolParms call; however, in AFP Version 1.1, only the Backup Date field may be set.
Backup Date (LONG)	new Backup Date

OUTPUTS: FPErr (LONG)

ERRORS:

ParamErr	unknown session refnum or volume identifier
BitmapErr	an attempt was made to set a parameter which cannot be set with this call; null bitmap

ALGORITHM: The Bitmap and Backup Date are passed to the server, which changes the date of the specified volume.

RIGHTS: The user must have previously called FPOpenVol for this volume.

PACKET FORMAT:

Command	
SetVolParms command	
Ø	
—	Volume ID
—	Bitmap
—	Backup Date

FPWrite

This call is used to write a block of data to an open fork.

INPUTS:	SRefNum (INT)	session refnum
	OForkRefnum (INT)	open fork refnum
	Offset (LONG)	byte offset from the beginning or end of the fork at which the write is to begin (should be negative to indicate a point within the fork relative to the end)
	ReqCount (LONG)	number of bytes to write
	Start/EndFlag (BIT)	one-bit flag indicating whether the Offset field is relative to the beginning or end of the fork (all other bits must be zero): 0 = relative to the beginning of the fork 1 = relative to the end of the fork

Fork data

OUTPUTS:	FPErrror (LONG)	
	ActCount (LONG)	number of bytes actually written to the fork
	LastWritten (LONG)	the number of the byte just past the last byte written

ERRORS:	ParamErr	unknown session refnum or open fork refnum
	AccessDenied	fork was not opened for Write
	LockErr	some or all of requested range is locked by another user
	DiskFull	no more space on the volume

ALGORITHM: The fork is written starting Offset bytes from the beginning or end of the fork. If the write extends beyond the end-of-fork, the fork is extended. If part of the range is locked by another user, a LockErr will be returned and no data will be written.

If multiple writers are concurrently modifying the fork, they may each have a different notion of the end-of-fork, although the server always knows the correct end-of-fork. It is for this reason that the "write relative to end-of-fork" feature is provided. The number of the byte just past the last byte written is returned, since the end-of-fork may have been different from the user's notion at the time at which the write was made.

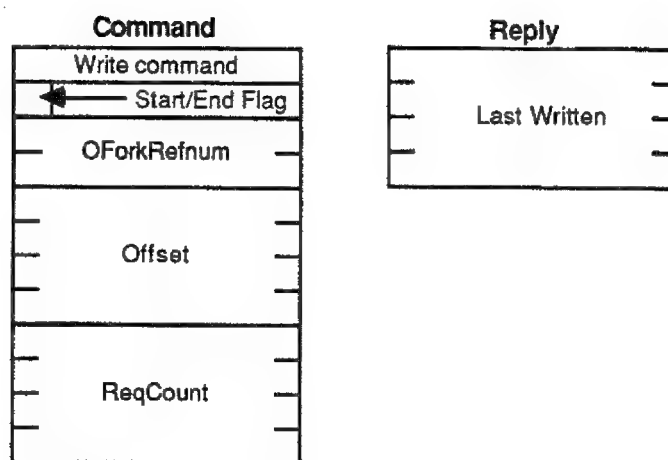
The fork data to be written is transmitted to the server in an intermediate exchange of Session Protocol packets.

The file's Mod Date is not changed until the fork is closed.

RIGHTS: The fork must have been opened for Write.

NOTES: Locking the range before writing to it is highly recommended, since the underlying transport mechanism may force the request to be broken up into multiple small requests. Although the range may not be locked when the call begins execution, it is possible for another session to lock some or all of the range before this call completes, causing the write to succeed partially.

PACKET FORMAT:



Chapter 8

AFP's Use of ASP

The AppleTalk Filing Protocol uses the transport services of the AppleTalk Session Protocol (ASP). This chapter discusses the manner in which the AFP calls are conveyed via ASP. For this purpose we refer to various terms defined in the ASP specification document. For these definitions, the reader must turn to that document.

In the case of all these calls, the AFP level variable *FPError* is returned as the ASP-level *CmdResult*. Also, the AFP-level Command packet is conveyed as an ASP-level *Command block*, while the AFP-level Reply packet is returned as the ASP-level *Command Reply block*.

Finding a Server

To find an AFP server, an NBP request must be issued for objects of type "AFPServer".

Getting Server Information

A workstation AFP client makes an *FPGetSrvrInfo* call to obtain server information needed before an AFP level login can be attempted. The *FPGetSrvrInfo* call is converted by the workstation's AFP into an ASP level *SPGetStatus* call. The server information is returned by ASP as the *SPGetStatus* call's *Command reply block*. Note that the *SAddr* (internet address of the file server's SLS) supplied with the *FPGetSrvrInfo* call is used as the *SLSEntityIdentifier* required by the *SPGetStatus* call.

Login On the File Server

Having obtained the server information, the AFP client initiates the process for logging in on the file server. This is achieved by AFP in a two step fashion.

First, AFP issues an *SPOpenSession* call to establish an ASP-level session. If for some reason this ASP session can not be established, then an error is returned to the workstation's AFP client and login is abandoned.

Once the ASP session has been opened, then the workstation's AFP issues an *SPCommand* call to send the *FPLogin's Command* to the file server's AFP. This *FPLogin Command* is sent as the *Command block* of the *SPCommand*. Beyond this point a series of *SPCommand's* may be necessary to complete the login process (this depends on the User Authentication Method being employed in the login).

If at any point, in this process the server returns an error indicating that the login can not be completed successfully, then the workstation's AFP must make an *SPCloseSession* call to close the ASP session. This is vital, otherwise the session will remain open for no purpose. After closing the session, the workstation's AFP returns an appropriate error message to its client.

Logout of the File Server

When the workstation client wishes to terminate its conversation with the file server, it issues an *FPLogout* call to its AFP. The workstation's AFP again uses a two step process to carry out this logout.

First, it issues an *SPCommand* to convey the *FPLogout* command to the server. Having done so, it then terminates the corresponding ASP session by making an *SPCloseSession* call.

Other AFP Calls

Every other AFP calls (with the exception of *FPWrite*) is conveyed by AFP via an *SPCommand* call.

An *FPWrite* call is sent by the workstation's AFP by making an *SPWrite* call.

It should be noted that in the case of the *FPRead*, *FPWrite*, and *FPEnumerate* calls, a partial completion of the call is possible. By this we mean that less than the desired length is read (or written). There are two reasons for this.

First, in the case of an *FPRead* call, the end-of-file may be reached before the requested number of bytes have been read. In the case of an *FPRead* or *FPWrite* call, a locked byte range may be reached before all the requested bytes have been read or written.

A second reason is related to the *QuantumSize* related to the session protocol. In the case of AppleTalk, the largest data block that may be written or read through ASP is equal to 4624 bytes, the *QuantumSize* size. Thus the underlying *SPCommand* (in the case of an *FPRead* or *FPEnumerate*) or *SPWrite* (in the case of an *FPWrite*) will complete with an actual received or written reply size smaller than the requested value. Therefore, if no error was returned, the AFP must issue an additional *SPCommand* or *SPWrite* call to complete the AFP request.

Although AFP may have to issue several ASP calls to complete a single AFP command, the first ASP command should convey the actual sizes requested by the user to allow the server a chance to optimize. Subsequent ASP calls should convey sizes that have been adjusted to reflect how much of the original command has already been completed.

Appendix A

User Authentication Methods

As noted in the main body of this document, AFP version 1.1 provides three standard user authentication methods. These correspond to UAM strings "NoUserAuthent", "Cleartxt passwd" and "Randnum exchange". [Note: these strings should be used in a case-insensitive fashion].

No User Authentication

The first of these in fact corresponds to no user authentication (UAM = "No User Authent") and as such needs no specification. Thus, no user name or password information is required in the *FPLogin* command which therefore has no *UserAuthInfo* field.

Any server can accept a user login using this method. If the server handles the user-authentication-based directory-level access control mechanism, then it must assign the user a special User ID and Group ID for that session, such that the user only obtains world's access rights for every directory in every server volume.

User Authentication with Clear Text Password Transmission

The second standard method employs the transmission of the user's password in clear text (in addition to the user's name) in the *FPLogin* command packet. The *UserAuthInfo* part of the *FPLogin* command consists of the user's name (a string of up to 31 characters which follows the UAM field in the packet without padding) followed by a possible null byte and then the user's password. The *user's password* is an 8-byte quantity. If the user provides a shorter password then it must be padded (suffixed) with null (\$00) bytes to make its length equal to 8 bytes. The permissible set of characters in passwords consists of all 8-bit characters with the most significant bit equal to 0.

User name comparison in servers must be case insensitive, but password comparison is to be case sensitive in this *UserAuthenticationMethod*. Of course, one could create a new *UserAuthenticationMethod* that performed case-insensitive password comparison.

It should be noted that this method should be used by workstations only if it is known that the intervening network is secure against "wire tapping", otherwise the password information can be picked up out of *FPLogin* command packets by anyone tapping the network.

User Authentication Based on Random Number Exchange

In environments where the network is not secure against tapping, a more secure method based on a random number exchange between the server and the workstation can be used. In this method, the user's password is never sent over the network and hence cannot be picked up by tapping. In fact, it is essentially impossible (as secure as the basic encryption method) to derive the password from the information that is sent over the network.

The underlying idea of the method is that the server knows the user's password and it wants to find out if the user trying to log in knows this password as well. In a sense, the method tries to determine if the user and the server "share the same secret", the password, without sending the secret information over the network.

First, the user sends the *FPLogin* command packet with UAM corresponding to "Randnum exchange" and the *UserAuthInfo* containing the user's name string, which follows the UAM in the packet without padding.

Upon receiving this packet, the server examines its user data base to determine if this is a valid user name. If the user name is not found in the user data base, then an error is sent to the workstation indicating this result and login is denied. If the name is found in the user data base, then the server generates a random number (64 bits in length) and sends it back to the workstation. This is returned as the reply to the *FPLogin* call, along with an ID number and an *AuthContinue* error. Although not really an "error", this value is returned to indicate that all is well so far but the user is not yet authenticated.

The workstation uses the NBS data encryption standard (DES) algorithm to encrypt the random number, using the user's password (case-sensitive, same format as in the clear text UAM) as the encryption key. It sends the encrypted value (64 bits) back to the server in the *UserAuthInfo* parameter of the *FPLoginCont* call along with the ID number returned from the *FPLogin* call. This ID number merely helps the server associate the two calls and is not used again. The server compares the workstation's encrypted value with the encrypted value obtained using the password from its user data base. If the two encrypted values match, then the user is considered to have been authenticated and the login is successful. The *FPLoginCont* call will return *NoErr* if so, *UserNotAuth* if not. In either case, no reply data is returned.

Appendix B

Long/Short Name Management Algorithms

Files and directories in AFP possess two names, a Long Name and a Short Name, because of differences in Macintosh and MS-DOS naming rules: Macintosh permits file and directory names to be made up of at most 32 characters, where valid characters may be any printable ASCII code except colon (\$3A). MS-DOS is more restrictive; names may be up to eight alphanumeric (plus other) characters, optionally followed by period (\$2E) and a one-to-three alphanumeric (plus other) character extension. The dual naming convention was devised to allow file servers to present to Macintosh and MS-DOS users a name space that is compatible with each of their respective file systems. Yet in order to ensure that an object can be uniquely specified by either name, some rules must be adhered to.

AFP naming rules are such that any MS-DOS name can be used directly as an object's Short Name, and any Macintosh name can be used directly as a Long Name. It is up to the file server to generate the "other" name for each object, deriving it from the first and making the name as "close" as possible to minimize semantic loss. Since Long Name format is a superset of Short Name format, some cases are easily handled. For example, if an MS-DOS machine generates a Short Name for an object, the object's Long Name can be made the same as the Short Name. Deriving a Short Name from a Long Name is not as simple, and AFP does not stipulate an exact algorithm for doing this, so that different servers may perform it differently. However, it is crucial that no two objects in a given directory have the same Short Name or the same Long Name. This ensures that any object can be uniquely specified by either name.

Below are the algorithms intended for use by file servers for creating and maintaining a consistent dual name space.

When an object is created (either by *FPCreateFile* or *FPCreateDir*), the caller supplies the object's name and a name type that indicates whether the name is of Short format or Long format. The server must of course first check the name itself to verify that it conforms to the specified format. The algorithm below describes how servers are to assign Short and Long names to objects:

```

IF name type is Short OR name is in Short format
THEN      check for new name in list of short names
            IF name already exists
            THEN      return ObjectExists error
            ELSE set object's Short and Long names to new name

ELSE { name type is Long OR name is in Long format }
    * check for new name in list of long names
    IF name already exists
    THEN      return ObjectExists error
    ELSE set object's Long Name to new name
           derive Short Name from Long Name
  
```

This algorithm is to be used for renaming as well (*FPRename*). When an object is renamed, both names will be changed using the above rules.

Note that this algorithm mandates that whenever an object is named with a Short Name format name, the Long Name will always match. This is done so that servers need only check one list, Long Names or Short Names, for each creation or renaming operation.

An interesting problem can arise. Consider the case in which a server contains a Macintosh file with a Long Name *MacFileWithLongName* and a Short Name *MacFile* which the server derived from the Long Name. This Short Name is not normally visible to Macintosh workstations. Now suppose that a user at a Macintosh decides to create a new file in the same directory with the Long Name *MacFile*. The call must fail, because by the above algorithm the new file's Short Name and Long Name must both be set to *MacFile* since it is in Short format. Hence the user is returned an *ObjectExists* error even though the directory does not appear to contain a file by that Long Name.

Appendix C

File Sharing Modes

AFP provides much functionality to control the sharing of files, primarily at the Creator/Group/World Access Rights level but also at the time of file fork open. To perform the latter function, the server must enforce the synchronization rules presented below.

When a fork of a file is opened, the user indicates what kind of Access Mode is needed: *Read*, *Write*, *Read/Write*, or *None* (the latter allows no further access to the fork except to close it, yet it is included since it may be useful as a synchronization primitive). In addition, the user indicates to the server a Deny Mode: just what rights should be denied to others trying to open the fork while the current users have it open. Users that subsequently try to open the fork could be denied *None*, *Read*, *Write*, or *Read/Write* access.

An *FPOpenFork* call could fail for several reasons:

- 1) The user may not possess the correct rights as Owner/Group/World to allow the desired access. An *AccessDenied* error is returned to the second user.
- 2) The fork may already be open with a Deny Mode that prohibits the second user's desired access. For example, the first user opened the fork with Deny Mode indicating *Deny Write*, and a second user tries to open the fork for *Write*. A *DenyConflict* error is returned to the second user.
- 3) The fork may already be open with a Deny Mode that conflicts with the second user's desired Deny Mode. For example, the first user opened the fork for *Write* and *Deny None*. The second user tries to open the fork with Deny Mode indicating *Deny Write*. This request cannot be granted since the fork is already open for *Write*. A *DenyConflict* error is returned to the second user.

Deny Modes are cumulative in that each successful open of a fork "combines" its Deny Mode with previous Deny Modes. That is, if the first open sets a Deny Mode of *Deny Read* and the second sets a Deny Mode of *Deny Write*, the fork's Current Deny Mode (CDM) will be *Deny Read/Write*. *Deny None* and *Deny Read* combine to form a Current Deny Mode of *Deny Read*. Likewise, Access Modes are cumulative: if the first open gets *Read* access and the second gets *Write* access, the Current Access Mode (CAM) is *Read/Write*.

The rules for allowing or denying multiple opens of a fork depend on the Current Deny Mode, the Current Access Mode, and the Deny and Access Modes being requested in a new *FPOpenFork* call, and are summarized in the table below. A check mark indicates that the new open is allowed, otherwise it is denied.

		New Open Attempt Deny Mode and Access Mode															
		Deny R/W				Deny Write				Deny Read				Deny None			
		-	R	RW	W	-	R	RW	W	-	R	RW	W	-	R	RW	W
Current Deny Mode and Current Access Mode	Deny R/W	-	√			√				√				√			
		R				√								√			
		RW												√			
		W								√				√			
	Deny Write	-	√	√		√	√			√	√			√	√		
		R				√	√							√	√		
		RW												√	√		
		W								√	√			√	√		
	Deny Read	-	√		√	√			√	√			√	√			√
		R				√			√					√			√
		RW												√			√
		W								√			√	√			√
	Deny None	-	√	√	√	√	√	√	√	√	√	√	√	√	√	√	√
		R				√	√	√	√					√	√	√	√
		RW												√	√	√	√
		W								√	√	√	√	√	√	√	√

Figure C-1: Rules for Allowing or Denying Multiple Opens

Appendix D

Values of Command and Error Codes

Command Codes

The command codes used in the command packets are listed below. Each command code is a 16 bit integer sent in the packet high byte first. The values are given here in decimal (base 10) form.

Command	Value
ByteRangeLock	1
CloseVol	2
CloseDir	3
CloseFork	4
CopyFile	5
CreateDir	6
CreateFile	7
Delete	8
Enumerate	9
Flush	10
FlushFork	11
GetForkParms	14
GetSrvrInfo	15
GetSrvrParms	16
GetVolParms	17
Login	18
LoginCont	19
Logout	20
MapID	21
MapName	22
Move	23
OpenVol	24
OpenDir	25
OpenFork	26
Read	27
Rename	28
SetDirParms	29
SetFileParms	30
SetForkParms	31
SetVolParms	32
Write	33
GetFileDirParms	34
SetFileDirParms	35
OpenDT	48
CloseDT	49
GetIcon	51

GetIconInfo	52
AddAPPL	53
RmvAPPL	54
GetAPPL	55
AddComment	56
RmvComment	57
GetComment	58
AddIcon	192

Error Codes

Each call returns an Error code which is a 4-byte integer. The various error values are listed below together with their mnemonic names (these are the names used in Chapter 7). The values given below are in hexadecimal and decimal (base 10) form.

Error Mnemonic	Hex Value	Decimal Value
NoErr	\$0	0
AccessDenied	\$FFFFEC78	-5000
AuthContinue	\$FFFFEC77	-5001
BadUAM	\$FFFFEC76	-5002
BadVersNum	\$FFFFEC75	-5003
BitmapErr	\$FFFFEC74	-5004
CantMove	\$FFFFEC73	-5005
DenyConflict	\$FFFFEC72	-5006
DirNotEmpty	\$FFFFEC71	-5007
DiskFull	\$FFFFEC70	-5008
EOFErr	\$FFFFEC6F	-5009
FileBusy	\$FFFFEC6E	-5010
FlatVol	\$FFFFEC6D	-5011
ItemNotFound	\$FFFFEC6C	-5012
LockErr	\$FFFFEC6B	-5013
MiscErr	\$FFFFEC6A	-5014
NoMoreLocks	\$FFFFEC69	-5015
NoServer	\$FFFFEC68	-5016
ObjectExists	\$FFFFEC67	-5017
ObjectNotFound	\$FFFFEC66	-5018
ParamErr	\$FFFFEC65	-5019
RangeNotLocked	\$FFFFEC64	-5020
RangeOverlap	\$FFFFEC63	-5021
SessClosed	\$FFFFEC62	-5022
UserNotAuth	\$FFFFEC61	-5023
CallNotSupported	\$FFFFEC60	-5024
ObjectTypeErr	\$FFFFEC5F	-5025
TooManyFilesOpen	\$FFFFEC5E	-5026
ServerGoingDown	\$FFFFEC5D	-5027
CantRename	\$FFFFEC5C	-5028
DirNotFound	\$FFFFEC5B	-5029
IconTypeError	\$FFFFEC5A	-5030

Appendix E

List of Abbreviations

AFI	AppleTalk filing interface
AFP	AppleTalk Filing Protocol
ALAP	AppleTalk Link Access Protocol
ASP	AppleTalk Session Protocol
ATP	AppleTalk Transaction Protocol
CAM	Current Access Mode
CDM	Current Deny Mode
DDP	Datagram Delivery Protocol
GMT	Greenwich Mean Time
HFS	Macintosh's hierarchical file system
MFS	the earlier Macintosh (flat) file system
NBP	Name Binding Protocol
NFI	Native filing interface
SLS	Session listening socket
UAM	User authentication method



ULISSE s.r.l.
Servizio Distribuzione Apple
Via Pacini, 22
20131 Milano
Tel. (02) 2367783-2367406

su licenza esclusiva
Apple Computer S.p.A.

AppleTalk Manager Update

IAT03

**APDA Macintosh AppleTalk Update
Release Notes
February 25, 1987**

The APDA Macintosh AppleTalk Update consists of one 400K diskette, a document entitled "AppleTalk Manager Update, APDA Draft," and these release notes. The disk contains updates and additions to the current Macintosh AppleTalk utilities. The document provides technical details of many of the changes. It is a "pre-release" of a chapter to be published in *Inside Macintosh* volume 5; future updates to it will appear there. These release notes are meant to summarize and supplement that document.

Disk Contents: The APDA Macintosh AppleTalk Update Disk includes the following:

- (1) Version 4.0 of the AppleTalk High-Level ("Pascal") Interface for MPW. This version is a maintenance release which fixes some known bugs.
- (2) Version 48 of the AppleTalk drivers. This version, which is the one used in the AppleShare server, provides many additional features over the version provided in the MacPlus ROM (version 40).
- (3) Version 1.1 of the new Extended Protocol Package (XPP) driver. This driver, used in the AppleShare workstation, implements the workstation side of the AppleTalk Session Protocol (ASP).
- (4) MPW Pascal and C interfaces to the XPP driver.

Each of these components is summarized in the following pages. Since the distribution disk is an MFS volume, you may wish to copy it to an HFS volume before installation.

Important: As stated in your APDA membership agreement, these utilities are provided for development purposes only; no license to distribute in any way is implied. Please contact Apple Software Licensing if you wish to distribute any of the enclosed material.

AppleTalk High-Level Interface Version 4.0

Version 4.0 of the AppleTalk "Pascal" Interface is a maintenance release. It replaces the current version, 3.4. There are three major changes to this version:

- (1) elimination of the use of location \$E0
- (2) elimination of the 'atpl' ('{Libraries}AppleTalk') resource
- (3) bug fixes.

Each of these changes is described below. Note that there is no added support for any of the new calls available in the new AppleTalk drivers (version 48). See the section on "future plans."

Use of location \$E0: Previous versions of the Pascal Interface used low memory location \$E0. This created a number of problems. The use of this location has been eliminated.

The 'atpl' resource: Previous versions of the Pascal Interface included a resource of type 'atpl' which had to be included in either the system file or the app's resource fork. This resource was shipped as the file {Libraries}AppleTalk under MPW 1.0. It always took up about 5K on disk and in memory, regardless of how much of it was needed. Questions arose as to how to release the resource when it was no longer needed. This resource was called by a small bit of glue linked into your application out of Interface.o.

The current release eliminates the use of this resource altogether. The glue in interface.o has been replaced with the actual code from the resource. Thus the actual code is now linked directly into your application. Since the linker only includes that code which is needed, in general only 1-2K will be added to the size of your application's code. Since 5K will be deleted by the elimination of the 'atpl' resource, there should be a net space savings, both on disk and in memory. Issues such as how to free the resource are also eliminated.

Bug fixes: A number of small bugs have been fixed. These include:

- (1) DDPReadCancel could crash if you tried to cancel a DDPRead that wasn't there
- (2) ATPReqCancel and ATPRspCancel now ignore the async flag and always execute synchronously. This is because the ABRecord passed in these calls is the ABRecord of the call to be cancelled, not a new ABRecord for the cancel call. Thus there is no way to indicate completion of the call itself if it were asynchronous.

Installation procedure: Included on the release disk, in folder "PI4.0" is an MPW script which replaces the current glue in {Libraries}Interface.o with the new routines. It is called "NewInterface". To execute it, just set the current directory to that folder (pi4.0) and type NewInterface. This will copy the routines from file "ABPackage.o" into {Libraries}Interface.o (note that executing NewInterface twice will result in a number of warnings). Alternatively, you could just include the file ABPackage.o on your link command line before Interface.o. This will cause a number of duplicate label warnings, however.

Following installation, all programs using the Pascal Interface should be re-linked, and the 'atpl' resource can be removed from all disks.

Future plans: We intend this to be the final release of the current-style interface. We will continue to support this for some time. However, in the long run, we intend to move to a more straightforward and simple interface design. This interface, which will also support all of the new AppleTalk functionality, will be a parameter-block style interface, much like the current low-level file system interface detailed in *Inside Mac*, volume 4. Additional functionality will be added only where necessary to support operations that can not be performed in a higher-level language (for instance socket listeners).

Developers will be free to continue to use the old style interface for MacPlus-style AppleTalk functionality. However we believe in the long run it will be advantageous to move to the new style interface.

We expect this new style interface to be initially available in 2-3 months. Most calls will be essentially the same format as the high-level XPP calls included with this release.

AppleTalk Drivers Version 48

Version 48 of the AppleTalk drivers contains significant enhancements over the current version (40 in the MacPlus ROMs, and 41 in system file 3.2). These enhancements include new calls and additional resources (e.g. all dynamic sockets can be opened simultaneously; more than one NBP request can be outstanding at any one time). These drivers are used by the AppleShare server. We envision these drivers to be used in place of the MacPlus ROM drivers in two circumstances: (1) where one or more of the new calls is required and (2) in servers where additional resources are needed. These new drivers are completely compatible with the MacPlus ROM drivers, and can be used in any situation where the ROM drivers were used previously.

The developer should be aware however that use of these drivers on a MacPlus or 512Ke will result in the loss of about 11K of system heap, since they completely replace the ROM drivers and since more memory is required to provide the additional resources. This may require the system heap to be grown, depending on the application.

On an original ROM 512K Mac, these new drivers can be placed in the system file and will load on the first `_Open` call. On a MacPlus or 512Ke, however, ROM drivers supersede those in the system file, so the new drivers must be loaded in some other method. This release provides a system file INIT resource, INIT 22, which, for MacPlus ROMs, opens the new drivers out of a file called "AppleTalk" in the system folder. If the file is not there, the INIT resource does nothing. Note that this INIT must be installed in the system file so that it runs before any application INIT's (INIT 31's) which might open AppleTalk and start up a request. Apple intends to ship this INIT as part of future system files; however, for now, the developer must install it himself.

These new drivers are distributed as five resources in the file "AppleTalk" within the folder "newMPPATP". These are DRVR 9 and 10, and NBPC 1 and 2. In addition, this file contains a resource of type 'mppc'. This is a new resource used to pass information to the INIT. Its format will be specified in the future. This resource, as provided, tells the INIT to open AppleTalk in server mode (i.e. pick a server node number, a process which takes about 10 seconds). If you wish to have AppleTalk opened in workstation mode, do not include this resource in the "AppleTalk" file.

The INIT 22 resource is also included in the "AppleTalk" file. To use the new AppleTalk drivers for development purposes, you can install this INIT in your system file and copy the file "AppleTalk" into your system folder, deleting the 'mppc' resource if desired.

For development purposes, installation of the INIT resource can be done using ResEdit or the equivalent. However, for shipping products, installation from the product disk onto the user's disk should be done using the Apple Installer. The current version of that installer is provided in the "newMPPATP" folder, along with the installation script "AppleTalk INIT Script". This script, which should be run off the product disk containing the "AppleTalk" file in the system folder, installs the INIT into the user's system file, copies the "AppleTalk" file into the system folder, and sets the boot blocks to a minimum of 58K (this is also the size used on the AppleShare workstation disk).

Assembly language equates for the new calls in these drivers are distributed in the file "nATalkEqu.a" in the same folder. This file also contains assembly equates for the new XPP driver.

XPP Version 1.1

The XPP (Extended Protocol Package) driver implements the workstation side of the AppleTalk Session Protocol (ASP), and a small portion of the AppleTalk Filing Protocol (AFP). It is installed in the system file and used in the same manner as the other AppleTalk drivers. XPP is automatically installed by the AppleShare workstation installer. Developers wishing to use XPP without AppleShare workstation software will have to install it in the system file themselves.

The XPP driver is resource DRVR 40, and is contained in the file XPP1.1 in the folder "XPP". The assembly language interface to this driver is included in the file "nAtalkEqu.a" in the folder "newMPPATP".

XPP High-Level Interface

A high-level, parameter block style interface to XPP is provided for both MPW Pascal and MPW C. It is distributed in three files in the "XPP" folder. The files "nAppleTalk.p" and "ncAppleTalk.h" are the Pascal and C interfaces, respectively, to the XPP driver. The file nAppleTalk.a.o is the object file containing the actual interface code. This file should be included on the link command line. Also included in this folder are two example MPW tools: "exOpenSession.p" and "exOpenSession.c". These tools simply issue an OpenSession call to XPP, as specified by their arguments.

AppleTalk Manager Update

APDA Draft

Apple Technical Publications Department

January 23, 1987

Copyright © 1987 Apple Computer, Inc. All rights reserved.

Table of Contents

1	Summary of Changes
1	Changes to the AppleTalk Manager
2	Picking a node address in the server range
2	Sending Packets to one's own node
3	AppleTalk Transaction Protocol Changes
3	Sending an ATP request through a specified socket
4	Aborting ATP SendRequests
4	Aborting Get Requests
5	Name Binding Protocol Changes
5	Multiple Concurrent NBP Requests
5	Variable Resources
6	New AppleTalk Protocols
8	Echo Protocol
8	AppleTalk Session Protocol
9	What ASP does
9	What ASP does not do
9	AppleTalk Filing Protocol
10	Extended Protocol Package Driver
12	Version
12	Error Reporting
13	AppleTalk Filing Protocol errors
14	.XPP driver functions overview
14	Using AppleTalk Name Binding Protocol
14	Opening and closing sessions
14	Session maintenance
14	Commands on an open session
15	Getting server status information
16	Attention mechanism
16	The attention routine
17	Calling the XPP Driver
17	Using XPP
17	Allocating memory
17	Opening the XPP Driver
18	Example
19	Open errors
19	Closing the XPP driver
19	Close errors
19	Session control block
19	How to access the XPP driver
19	Using Pascal
20	General
23	AppleTalk Session Protocol
23	Note on result codes
28	AFP Implementation
28	Mapping AFP commands
29	AFPCall function
35	CCB sizes
36	XPP Driver Results Codes
37	Summary

APPLETALK MANAGER UPDATE

The AppleTalk Manager has been enhanced through the implementation of new protocols and increased functionality of the existing interface. This chapter describes these enhancements beginning with a brief summary of the changes that have been made. The remainder of the chapter provides detailed information about these changes.

Summary of Changes

The AppleTalk Manager provides services that allow Macintosh programs to interact with clients in devices connected to an AppleTalk network. The following is a brief summary of the changes that have been made to the AppleTalk Manager interface.

- At open time, the .MPP driver can be told to pick a node number in the server range. This is a more time consuming but more thorough operation than selecting a node number in the workstation range, and is required for devices acting as servers.
- Multiple concurrent NBP requests are now supported (just as multiple concurrent ATP requests have been supported). The KillNBP command has been implemented to abort an outstanding NBP request.
- ATP requests can now be sent *through* client-specified sockets, instead of having ATP pick the socket itself.
- The ability to send packets to one's own node is supported (although this functionality is, in the default case, disabled).
- Two new ATP abort calls have been added: KillSendReq and KillGetReq. KillSendReq is functionally equivalent to RelTCB, although its arguments are different. KillGetReq is a new call for aborting outstanding GetRequests.
- Additional machine-dependent resources have been added to support , for example, more dynamic sockets and more concurrent ATP requests.
- A new protocol called the Echo Protocol (EP) is supported.
- A new driver, .XPP has been added. The .XPP driver implements the workstation side of the AppleTalk Session Protocol (ASP), and a small portion of the AppleTalk Filing Protocol.

To determine if you are running on a machine that supports these enhanced features, check the version number of the .MPP driver (at offset DCtlQueue+1 in the Device Control Entry). A version number of 48 (NCVersion) or greater indicates the presence of the new drivers.

Note: With the exception of the XPP calls, there is currently no high-level support for any of the new calls. They can however, be accessed through appropriate PBControl calls.

CHANGES TO THE APPLETALK MANAGER

Changes to the AppleTalk manager increase functionality and resources. Picking a node address in the server range, sending packets to one's own node, multiple concurrent NBP requests, sending ATP requests through a specified socket and two new ATP calls are discussed in this section.

Picking a node address in the server range

Normally upon opening, the node number picked by the AppleTalk manager will be in the node number range (\$01-\$7F). It is possible to indicate that a node number in the server range (\$80-\$FE) is desired. Picking a number in the server range is a more time consuming but more thorough process and is required for server nodes, since it greatly decreases the possibility of a node number conflict.

To open AppleTalk with a server node number, an extended open call is used. An extended open call is indicated by having the immediate bit set in the Open trap itself. In the extended open call, bit 7 of the extension longword field (ioMix) indicates whether a server or workstation node number should be picked. Set this bit to 1 to request a server node number. The rest of this field should be zero, as should all other unused fields in the queue element. A server node number can only be requested on the first Open call to the .MPP driver.

Sending packets to one's own node

Upon opening, the ability to send a packet to one's own node (intra-node delivery) is disabled. This feature of the AppleTalk Manager can be manipulated through the SetSelfSend function. Once enabled, it is possible, at all levels, to send packets to entities within one's own node. An example of where this might be desirable is an application sending data to a print spooler which is actually running in the background on the same node.

Note that enabling (or disabling) this feature affects the entire node and should be performed with care. For instance, a desk accessory may not expect to receive names from within its own node as a response to an NBP lookup; enabling this feature from an application could break the desk accessory. All future programs should be written with this feature in mind.

SetSelfSend function

Parameter Block

--->	26	csCodeword	; always SetSelfSend
--->	28	NewSelfFlag byte	; new SelfSend flag
<---	29	OldSelfFlag byte	; old SelfSend flag

SetSelfSend enables or disables the intra-node delivery feature of the AppleTalk Manager. If NewSelfFlag is non-zero, the feature will be enabled; otherwise it will be disabled. The previous value of the flag will be returned in OldSelfFlag.

Result Codes	noErr	No error
--------------	-------	----------

AppleTalk Transaction Protocol Changes

Sending an ATP request through a specified socket

ATP requests can now be sent through client-specified sockets. ATP previously would open a dynamic socket, send the request through it, and close the socket when the request completed. The client can now choose to send a request through an already-opened socket. This also allows more than one request to be sent per socket. A new call, `NSendRequest`, has been added for this purpose. The function of the old `SendRequest` call itself remains unchanged.

NSendRequest function

Parameter block

--->	18	userData	longword	;user bytes
<---	22	reqTID	word	;transaction ID used in request
--->	26	csCode	word	;always sendRequest
--->	28	atpSocket	byte	;socket to send request on
<-->	29	atpFlags	byte	;control information
--->	30	addrBlock	longword	;destination socket address
--->	34	reqLength	word	;request size in bytes
--->	36	reqPointer	pointer	;pointer to request data
--->	40	bdsPointer	pointer	;pointer to response BDS
--->	44	numOfBufs	byte	;number of responses expected
--->	45	timeOutVal	byte	;timeout interval
<---	46	numOfResps	byte	;number of responses received
<-->	47	retryCount	byte	;number of retries
<---	48	intBuff	word	;used internally

The `NSendRequest` call is functionally equivalent to the `SendRequest` call, however `NSendRequest` allows you to specify, in the `atpSocket` field, the socket *through* which the request is to be sent. This socket must have been previously opened through an `OpenATPSkt` request (otherwise a `badATPSkt` error will be returned). Note that `NSendRequest` requires two additional bytes of memory at the end of the parameter block, immediately following the `retryCount`. These bytes are for the internal use of the AppleTalk Manager and should not be modified while the `NSendRequest` call is active.

There is a machine-dependent limit as to the number of concurrent `NSendRequests` that can be active on a given socket. If this limit is exceeded, the error `tooManyReqs` will be returned.

Note that one additional difference between `SendRequest` and `NSendRequest` is that an `NSendRequest` can only be aborted by a `KillSendReq` call (see below), whereas a `SendRequest` can be aborted by either a `RelTCB` or `KillSendReq` call.

Result Codes	noErr	No error
	reqFailed	Retry count exceeded
	tooManyReqs	Too many concurrent requests
	noDataArea	Too many outstanding ATP calls
	reqAborted	Request cancelled by user

Aborting ATP SendRequests

The RelTCB call is still supported, but only for aborting SendRequests. To abort NSendRequests, a new call, KillSendReq has been added. This call will abort both SendRequests and NSendRequests. KillSendReq's only argument is the queue element pointer of the request to be aborted. The queue element pointer is passed at the offset of the KillSendReq queue element specified by AKillQE1.

KillSendReq function

Parameter block

---	26	csCode	word	; always KillSendReq
---	44	AKillQE1	pointer	; pointer to queue element

KillSendReq is functionally equivalent to RelTCB, except that it takes different arguments and will abort both SendRequests and NSendRequests. To abort one of these calls, place a pointer to the queue element of the call to abort in AKillQE1 and issue the KillSendReq call.

Result Codes	noErr	No error
	cbNotFound	AKillQE1 does not point to a SendReq or NSendReq queue element

Aborting Get Requests

ATP GetRequests can now be aborted through the KillGetReq call. This call looks and works just like the KillSendReq call, and is used to abort a specific GetRequest call. Previously it was necessary to close the socket to abort all GetRequest calls on the socket.

KillGetReq function

Parameter block

---	26	csCode	word	; always KillGetReq
---	44	AKillQE1	pointer	; pointer to queue element

KillGetReq will abort a specific outstanding GetRequest call (as opposed to closing the socket which aborts all outstanding GetRequests on that socket). The call will be completed with a reqAborted error. To abort a GetRequest, place a pointer to the queue element of the call to abort in AKillQE1 and issue the KillGetReq call.

Result Codes	noErr	No error
	cbNotFound	AKillQE1 does not point to a GetReq queue element

Name Binding Protocol Changes

Multiple Concurrent NBP Requests

NBP now supports multiple concurrent active requests. Specifically, a number of LookupNames, RegisterNames and ConfirmNames can all be active concurrently. The maximum number of concurrent requests is machine dependent; if it is exceeded the error tooManyReqs will be returned. Active requests can be aborted by the KillNBP call.

KillNBP function

Parameter block

->	26	csCode	word	; always KillNBP
->	28	NKillQE1	pointer	; pointer to queue element

KillNBP is used to abort an outstanding LookupName, RegisterName or ConfirmName request. To abort one of these calls, place a pointer to the queue element of the call to abort in NKillQE1 and issue the KillNBP call. The call will be completed with a ReqAborted error.

Result Codes	noErr	No error
	cbNotFound	NKillQE1 does not point to a valid NBP queue element

Variable resources

All dynamic sockets (\$80 through \$FE) can now be opened concurrently in addition to twelve (12) static sockets.

The following resources have also been increased:

- Number of protocols handlers (MPP)
- Number of concurrent SendRequests (ATP)
- Number of ATP sockets
- Number of concurrent XO Send Responses (ATP)
- Number of data areas (ATP)
- Number of concurrent NBP requests

NEW APPLETALK PROTOCOLS

The following protocols have been added to the AppleTalk Manager:

- Echo Protocol
- AppleTalk Session Protocol (workstation side)
- AppleTalk Filing Protocol (small portion of the workstation side)

The AppleTalk system architecture consists of a number of protocols arranged in layers. Each protocol in a specific layer provides services to higher-level layers (knowns as the protocol's clients) by building on the services provided by the lower-level layers. Figure 1 shows the AppleTalk Protocols and their corresponding network layers.

In Figure 1, the lines indicate the interaction between the protocols. Note that like the Routing Table Maintenance Protocol, the Echo Protocol is not directly accessible to Macintosh programs.

The details of these protocols are provided in *Inside AppleTalk*.

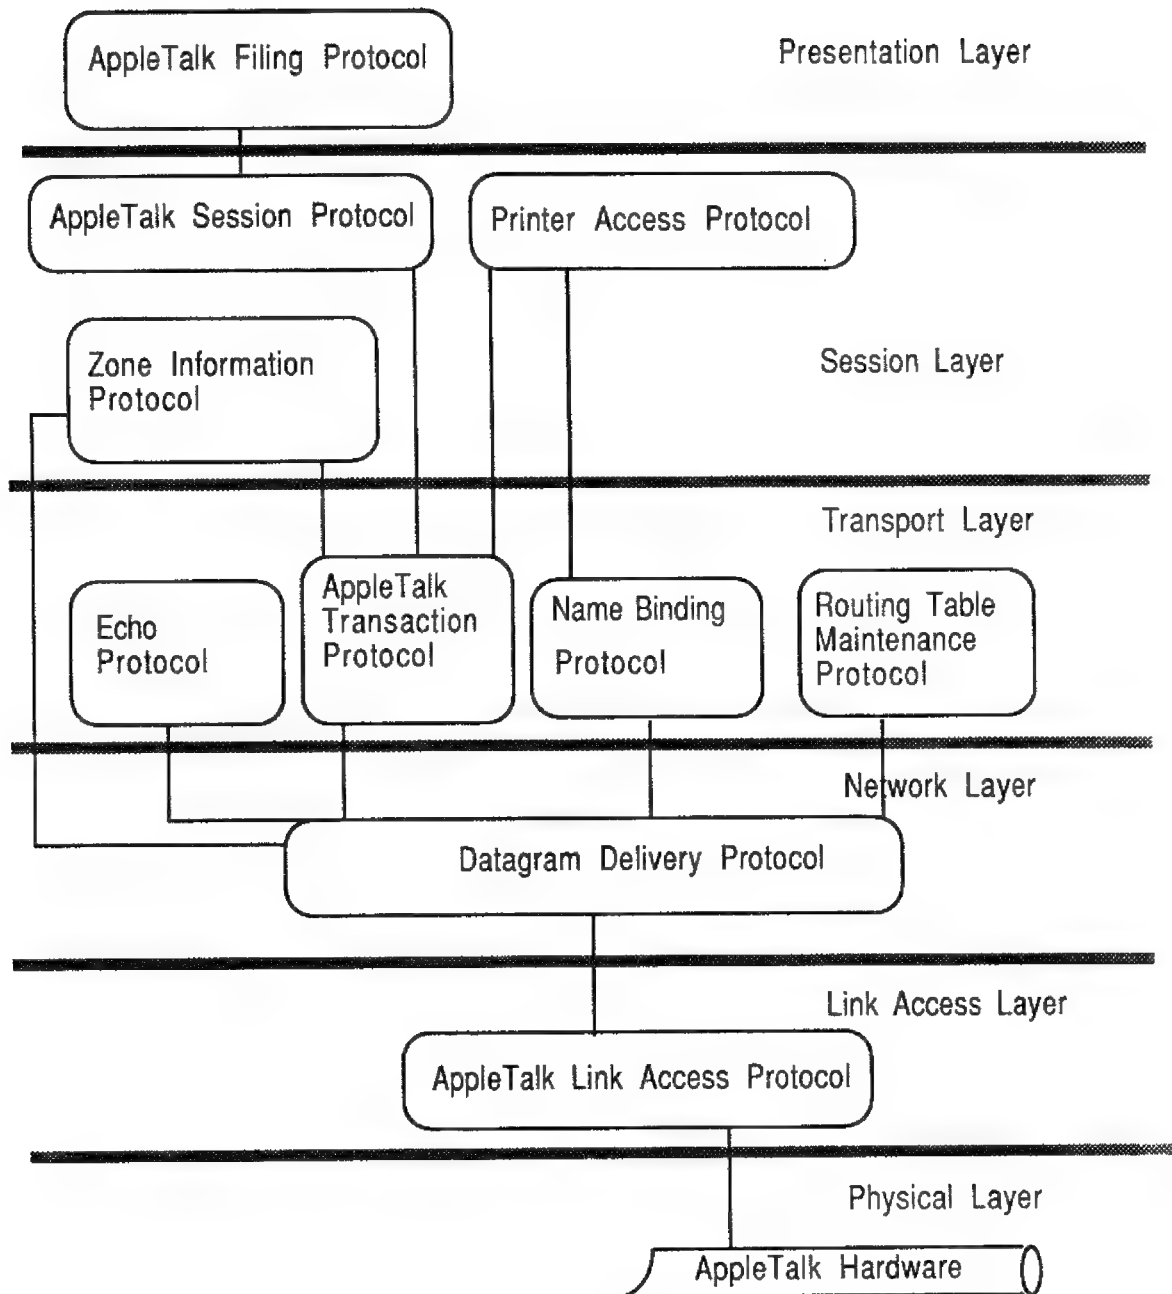


Figure 1. AppleTalk Protocols and OSI Network Layers

Echo Protocol

The **Echo Protocol** (EP) provides an echoing service through static socket number 4 known as the **echoer socket**. The **echoer** listens for packets received through this socket. Any correctly formed packet sent to the echoer socket on a node, will be echoed back to its sender.

This simple protocol can be used for two important purposes:

1. EP can be used by any Datagram Delivery Protocol (DDP) client to determine if a particular node, (known to have an echoer) is accessible over an internet.
2. EP is useful in determining the average time it takes for a packet to travel to a remote node and back. This is very useful in developing client-dependent heuristics for estimating the timeouts to be specified by clients of ATP, ASP and other protocols.

Programs can not access EP directly via the AppleTalk Manager. The EP implementation exists solely to respond to EP requests sent by other nodes. EP is a DDP client residing on statically-assigned socket 4, the echoing socket. Clients wishing to send EP requests (and receive EP responses) should use the Datagram Delivery Protocol (DDP) to send the appropriate packet.

AppleTalk Session Protocol

The **AppleTalk Session Protocol** (ASP) provides for the setting up, maintaining and closing down of a **session**. A session is a logical relationship between two network entities, a **workstation** and a **server**. The workstation tells the server what to do and the server responds with the appropriate actions. ASP makes sure that the sessions dialog is maintained in the correct sequence and that both ends of the conversation are properly participating.

ASP will generally be used between two communicating network entities where one is providing a service to the other (i.e. a server is providing a service to a workstation) and the service provided is *state-dependent*. That is, the response to a particular request from an entity is dependent upon other previous requests from that entity. For example, a request to read bytes from a file is dependent upon a previous request to open that file in the first place. However, a request to return the time of day is independent of all such previous requests.

When the service provided is state-dependent, requests must be delivered to the server in the same order as generated by the workstation. ASP guarantees requests are delivered to the server in the order in which they are issued, and that duplicate requests are never delivered (another requirement of state-dependent service).

What ASP does

ASP is an asymmetric protocol, providing one set of services to the workstation and a different set of services to the server.

ASP workstation clients initiate (open) sessions, send requests (commands) on that session and close sessions down. ASP server clients receive and respond (through command replies) to these requests. ASP guarantees that these requests are delivered in the same order as they are made, and without duplication. ASP is also responsible for closing down the session if one end dies or becomes unreachable and will inform its client (either server or workstation) of the action.

ASP also provides various additional services, such as allowing a workstation to obtain server status information without opening a session to a server, writing blocks of data from the workstation to the server end of the session, and the ability for a server to send an attention message to the workstation.

ASP assumes that the workstation client has a mechanism for looking up the network address of the server with which it wants to set up a session (generally this is done using the AppleTalk Name Binding Protocol).

Both ends of the session periodically check to see that the other end of the session is still responsive. If one end dies or becomes unreachable the other end closes the session.

ASP is a client of ATP and calls ATP for transport services.

What ASP does not do

ASP does not:

- allow the server to send commands to the workstation. The server is allowed to alert the workstation through the server's *attention* mechanism only.
- understand or interpret the syntax or the semantics of the commands sent to the server by the workstation.
- provide a user authentication (password) mechanism.
- insure that consecutive commands *complete* in the order in which they were sent (and delivered) to the server.

Note: The XPP driver does implement the workstation side of the AppleTalk Filing Protocol login command.

AppleTalk Filing Protocol

The AppleTalk Filing Protocol (AFP) allows a workstation on an AppleTalk network to access files on an AFP file server. AFP specifies a complex remote filing system containing user authentication and an access control mechanism that supports volume and folder-level access rights. For details of AFP, refer to the *AFP Draft Proposal*.

EXTENDED PROTOCOL PACKAGE DRIVER

The Extended Protocol Package (XPP) driver is intended to implement several AppleTalk communication protocols in the same package for ease of use. The .XPP driver currently consists of two modules that operate on two levels: the low-level implements the workstation side of AppleTalk Session Protocol, and the high-level implements (a small portion of) the workstation side of the AppleTalk Filing Protocol.

This driver adds functionality to the AppleTalk manager by providing services additional to those provided in the .MPP and .ATP drivers. Figure 2 shows the Macintosh AppleTalk drivers and the protocols accessible through each driver.

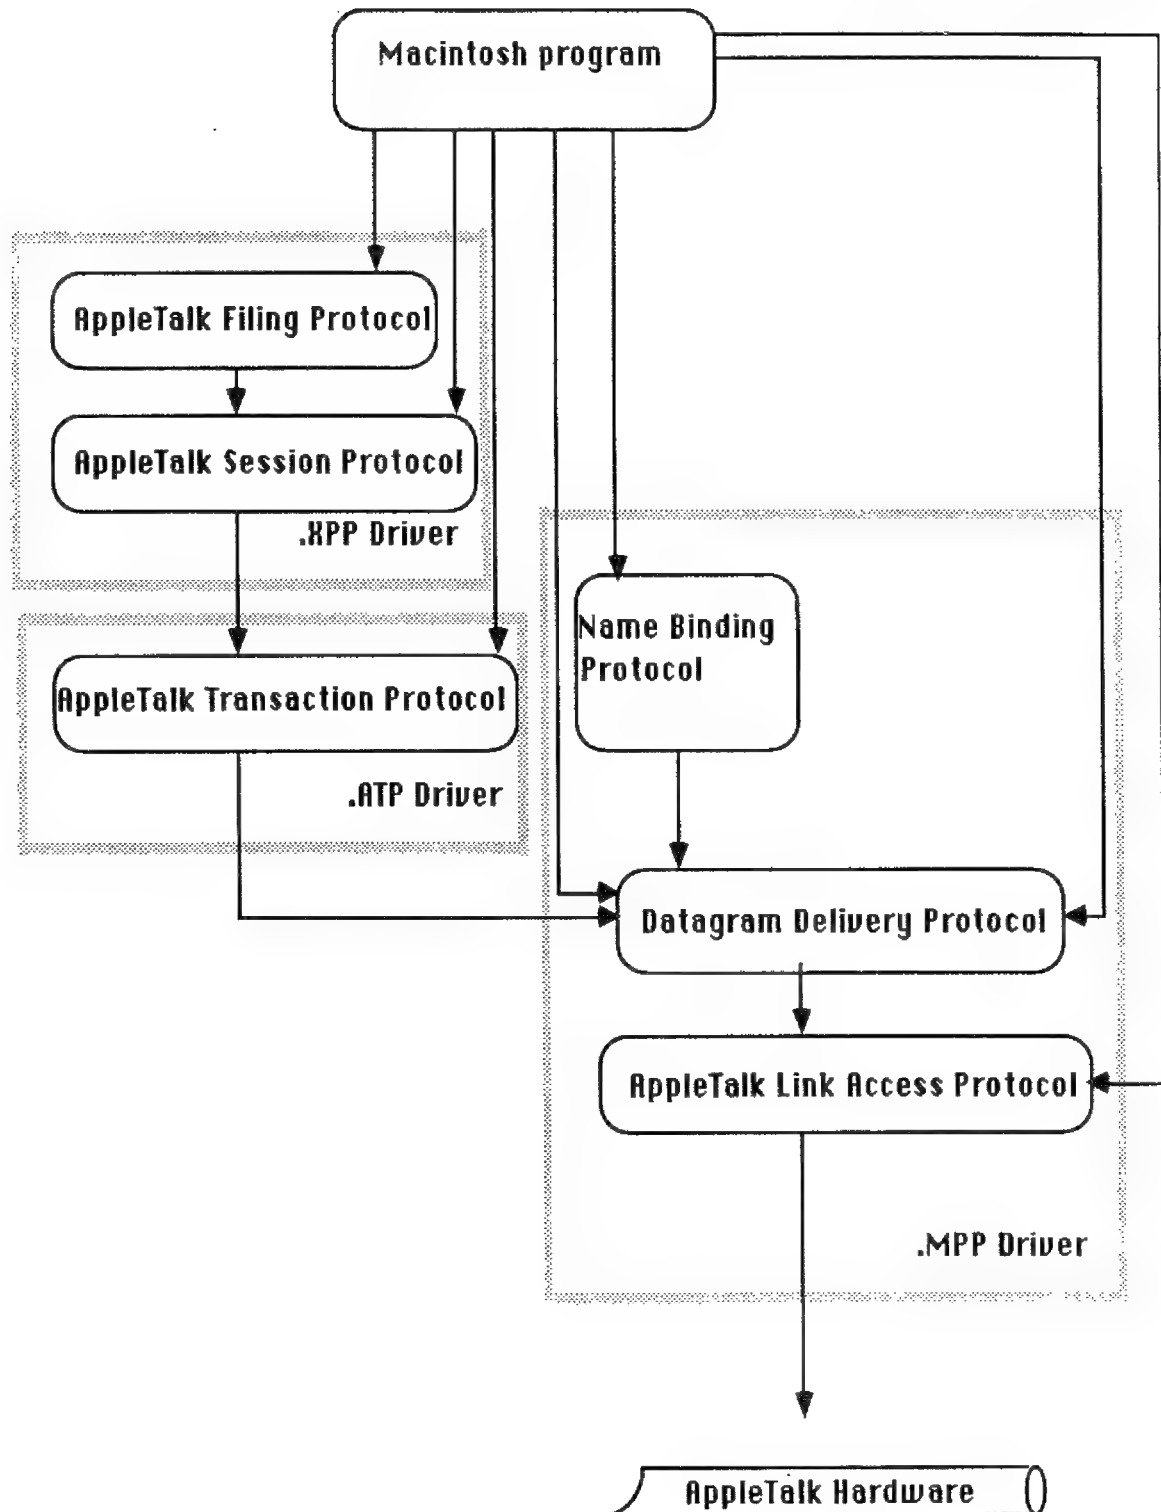


Figure 2. Macintosh AppleTalk Drivers

The .XPP driver maps an AFP call from the client workstation into one or more ASP calls. XPP provides one client-level call for AFP.

The implementation of AFP in the .XPP driver is very limited. Most calls are a very simple one-to-one mapping from an AFP call to an ASP command without any interpretation of the syntax of the AFP command by the .XPP driver. Refer to the *Mapping AFP commands* section of this document for further information.

Version

The .XPP driver supports ASP Version (hex) \$100, as described in *Inside AppleTalk*.

Error reporting

Errors are returned by the .XPP driver in the ioResult field of the Device Manager Control calls.

The error conditions reported by the .XPP driver may represent the unsuccessful completion of a routine in more than just one process involved in the interaction of the session. System-level, .XPP driver, AppleTalk, and server errors can all turn up in the ioResult field. Note that an ASP server error actually results from the activity of the server end of the transaction but is reported through the .XPP driver.

AFP calls also return codes indicating the unsuccessful completion of AFP commands in the Command Result field of the parameter block (described below).

An application using the .XPP driver should respond appropriately to error conditions reported from the different parts of the interaction. As shown in figure 3, the following errors can be returned in the ioResult field:

1. System-level errors

System errors returned by the .XPP driver indicate such conditions as the driver not being open or a specific system call not being supported. For a complete list of result codes returned by the Macintosh system software, refer to *Inside Macintosh*, Appendix A.

2. XPP errors (for example, *session not opened*)

The .XPP driver can also return errors resulting from its own activity (for example, the referenced session isn't open). The possible .XPP driver errors returned are listed in the .XPP driver results codes section with each function that can return the code.


3. AppleTalk Errors (returned from lower-level protocols)

XPP may also return errors from lower-level protocol, (for example, socket not open).

Possible error conditions and codes are described in *Inside Macintosh*, Volume 2, Chapter 10, "The AppleTalk Manager".

4. An ASP-specific error could be returned from an ASP server in response to a failed *OpenSession* call. Errors of this type, returned by the server to the workstation, are

documented both in *Inside AppleTalk* Section 11, "AppleTalk Session Protocol" and in the .XPP driver results code section of this document.



Error Types

1. System Error
2. XPP Error
3. AppleTalk Error
4. ASP Server Error
5. AFP Server Error

Figure 3. Error Reporting

AppleTalk Filing Protocol errors

In addition, the AppleTalk Filing Protocol defines errors that are returned from the server to the workstation client. These errors are returned in the the *CmdResult* field of the parameter block (error type 5 in Figure 3). This field is valid if no system-level error is returned by the call. Note that at the ASP level, the *CmdResult* field is client-defined data and may not be an error code.

.XPP driver functions overview

The paragraphs below describe the implementation of ASP in the .XPP driver. For more detailed information about ASP, refer to *Inside AppleTalk*, Section 11, "AppleTalk Session Protocol (ASP)."

Using AppleTalk Name Binding Protocol

A server wishing to advertise its service on the AppleTalk network calls ATP to open an ATP responding socket known as the *session listening socket* (SLS). The server then calls the Name Binding Protocol (NBP) to register a name on this socket. At this point, the server calls the server side of ASP to pass it the address of the SLS. Then, the server starts listening on the SLS for session opening requests coming over the network.

Opening and closing sessions

When a workstation wishes to access a server, the workstation must call NBP to discover the SLS for that server. Then the workstation calls ASP to open a session with that server.

After determining the SLS (address) of the server, the workstation client issues an `OpenSession` (or `AFPLLogin`) call to open a session with that server. As a result of this call, ASP sends a special `OpenSession` packet (an ATP request) to the SLS; this packet carries the address of a workstation socket for use in the session. This socket is referred to as the workstation session socket (WSS). If the server is unable to set up the session, it returns an error. If the request is successful, the server returns no error and the session is opened. The open session packet also contains a version number so that both ends can verify that they are speaking the same version of ASP.

The `AbortOS` function can be used to abort an outstanding `OpenSession` request before it has completed.

The workstation client closes the session by issuing a `CloseSession` (or `AFPLLogout`). The `CloseSession` call aborts any calls that are active on the session and closes the session. The session can also be closed by the server or by ASP itself, such as when one end of the session dies. The `CloseAll` call (which should be used with care) aborts every session that the driver has active.

Session maintenance

A session will remain open until it is explicitly terminated by the ASP client at either end or until one of the sessions ends dies or becomes unreachable.

Commands on an open session

Once a session has been opened, the workstation client can send a sequence of commands over the session to the server end. The commands are delivered in the same order as they are issued from the workstation end, and replies to the commands are returned to the workstation end.

Three types of commands can be made on an open session. These commands are `UserCommand`, `UserWrite`, and `AFPCall` functions described in the following paragraphs.

UserCommand calls are similar to ATP requests. The workstation client sends a command (included in a variable size command block) to the server client requesting it to perform a particular function and send back a variable size command reply. Examples of such commands vary from a request to open a particular file on a file server, to reading a certain range of bytes from a device. In the first case, a small amount of reply data is returned, in the second case a multiple-packet reply might be generated.

The .XPP driver does not interpret the command block or in any way participate in the command's function. It simply conveys the command block, included in a higher-level format, to the server end of the session, and returns the command reply to the workstation-end client. The command reply consists of a four-byte command result and a variable size command reply block.

UserWrite allows the workstation to convey blocks of data to the server. UserWrite is used to transfer a variable size block of data to the server end of the session and to receive a reply.

The AFPCall function provides a mechanism for passing an AFP command to the server end of an open session and receiving a reply. The first byte of the AFPCall command buffer contains the code for the AFP command that is to be passed to the server for execution. Most AFP calls are implemented through a very simple one-to-one mapping that takes the call and makes an ASP command out of it.

The AFPCall function can have one of four different formats. These four formats, which are basically all very similar, are described in detail below.

Getting server status information

ASP provides a service to allow its workstation clients to obtain a block of service status information from a server without the need for opening a session. The *GetStatus* function returns a status block from the server identified by the indicated address. ASP does not impose any structure on the status block. This structure is defined by the protocol above ASP.

Attention mechanism

Attentions are defined in ASP as a way for the server to alert the workstation of some event or critical piece of information. The ASP OpenSession and AFP login calls include a pointer to an attention routine in their parameter blocks. This attention routine is called by the .XPP driver when it receives an attention from the server and also when the session is closing as described below.

In addition, upon receiving an OpenSession call, or AFPLlogin call, the .XPP driver sets the first two bytes of the session control block (SCB) to zero. When the .XPP driver receives an attention, the first two bytes of the SCB are set to the attention bytes from the packet (which are always non-zero).

Note: A higher-level language such as Pascal may not wish to have a low-level attention routine called. A Pascal program can poll the attention bytes, and if they are ever nonzero, the program will know that an attention has come in (it would then set the attention bytes back to zero). Of course, two or more attentions could be received between successive polls and only the last one will be recorded.

The .XPP driver also calls the attention routine when the session is closed by either the server, workstation, or ASP itself (that is, timeout). In these cases, the attention bytes in the SCB are unchanged.

The attention routine

The attention routine is called at interrupt level and must observe interrupt conventions. Specifically, the interrupt routine can change registers A0 through A3 and D0 through D3 and it must not make any Memory Manager calls.

It will be called with

- D0 (word) equal to the SessRefnum for that session (see OpenSession Function)
- D1 (word) equal to the attention bytes passed by the server (or zero if the session is closing)

Return with an RTS (return from subroutine) to resume normal execution.

The next section describes the calls that can be made to the .XPP driver.

CALLING THE XPP DRIVER

This section describes how to use the XPP driver and how to call the XPP driver routines from assembly language and Pascal.

Using XPP

The XPP driver implements the workstation side of ASP and provides a mechanism for the workstation to send AppleTalk Filing Protocol (AFP) commands to the server.

Allocating memory

Every call to the XPP driver requires the caller to pass in whatever memory is needed by the driver for the call, generally at the end of the queue element. When a session is opened, the memory required for maintenance of that session (i.e., the Session Control Block) is also passed in.

For standard Device Manager calls, a queue element of a specific size equal to IOQELSize is allocated. When issuing many calls to XPP, it is the caller's responsibility to allocate a queue element that is large enough to accommodate the XPP driver's requirements for executing that call, as defined below. Once allocated, that memory can't be modified until the call completes.

Opening the XPP driver

To open the XPP driver, issue a Device Manager Open call. Refer to *Inside Macintosh*, Volume 2, Chapter 6, "The Device Manager." The name of the XPP driver is '.XPP'. Note that original Macintosh ROMs require that XPP be opened only once. With new ROMs, the XPP unit number can always be obtained through an *Open* call. With old ROMs only, the XPP unit number must be hard coded to XPPUnitNum (40) since only one open call can be issued to the driver.

The XPP driver cannot be opened unless AppleTalk is open. The application must ensure that the MPP and ATP drivers are opened, as described in *Inside Macintosh* Volume 2, pages 304-305.

The XPPLoaded bit (bit 5) in the PortBUse byte in low memory indicates whether or not the .XPP driver is open.

Example

The following is an example of the procedure an application might use to open the .XPP driver.

```

;-----
;
;   Routine: OpenXPP
;
;       Open the .XPP driver and return the driver refNum for it.
;
;       Exit:  D0 = error code (ccr's set)
;              D1 = XPP driver refNum (if no errors)
;
;       All other registers preserved
;
xppUnitNum    EQU    40                ; default XPP driver number
xppTfRNum     EQU    -(xppUnitNum+1)   ; default XPP driver refNum

OpenXPP
    MOVE.L A0-A1/D2,-(SP)              ; save registers
    MOVE    ROM85,D0                  ; Check ROM type byte
    BPL.S   @10                       ; branch if >=128K ROMs
    BTST    #xppLoadedBit,PortBUse; is the XPP driver open already?
    BEQ.S   @10                       ; if not open, then branch to Open code.
    MOVE    #xppTfRNum,D1             ; else use this as driver refnum.
    MOVEQ   #0,D0                     ; set noErr.
    BRA.S   @90                       ; and exit.

;
; XPP driver not open. Make an _Open call to it. If using a 128K ROM
; machine and the driver is already open, we will make another Open call to
; it just so we get the correct driver refNum.
;
@10          SUB    #ioQE1Size,SP      ; allocate temporary param block.
            MOVE.L SP,A0              ; A0 -> param block.
            LEA     XPPName, A1       ; A1 -> XPP (ASP/AFP) driver name.
            MOVE.L A1,ioFileName(A0) ; driver name into param block.
            CLR.B   ioPermssn(A0)     ; clear permissions byte
            _Open
            MOVE    ioRefNum(A0),D1   ; D1=driver refNum (invalid if error)
            ADD     #ioQE1Size,SP     ; deallocate temp param block.
@90          MOVE.L (SP)+,A0-A1/D2    ; restore registers
            TST     D0                ; error? (set ccr's)
            RTS

XPPName      DC.B   4                 ; length of string.
            DC.B   '.XPP'            ; driver name.

```


From Pascal, XPP can be opened through the OpenXPP call, which returns the driver's reference number.

```
FUNCTION OpenXPP (VAR XPPRefnum: INTEGER): OSErr;
```

Open errors

- Errors returned by System
- portInUse is returned if the AppleTalk port is in use by a driver other than AppleTalk or if AppleTalk is not open.

Closing the XPP driver

To close the XPP driver, call the Device Manager Close routine.

Caution: There is generally no reason to close the driver. Use this call sparingly, if at all. This call should generally be used only by system level applications.

Close Errors

- Errors returned by System
- Closerr (new ROMs only) returned if you try to close the driver and there are sessions active through that driver. When sessions are active, closerr is returned and the driver remains open.
- On old ROMs the driver is closed whether or not sessions are active and no error is returned. Results are unpredictable if sessions are still active.

Session Control Block

The session control block (SCB) is a non-relocatable block of data passed by the caller to XPP upon session opening. XPP reserves this block for use in maintaining an open session. The SCB size is defined by the constant SCBMemSize. The SCB is a locked block and as long as the session is open, the SCB can not be modified in any way by the application. There is one SCB for each open session. This block can be reused once a CloseSess call is issued and completed for that session or when the session is indicated as closed.

How to access the XPP driver

This section contains information for programmers using Pascal and assembly language routines.

Most XPP driver routines can be executed either synchronously (meaning that the application can't continue until the routine is completed) or asynchronously (meaning that the application is free to perform other tasks while the routine is executing).

Using Pascal

On asynchronous calls, the caller may pass a completion routine pointer in the parameter block, at offset IOCompletion. This routine will be executed upon completion of the call. It is executed at interrupt level and must not make any memory manager calls. If it uses

AppleTalk Manager Update

application globals, it must ensure that register A5 is set up correctly; for details see SetupA5 and RestoreA5 in the Operating System Utilities chapter of *Inside Macintosh*. If no completion routine is desired, IOCompletion should be set to NIL.

Asynchronous calls return control to the caller with result code of noErr as soon as they are queued to the driver. This isn't an indication of successful completion. To determine when the call is actually completed, if you don't want to use a completion routine, you can poll the ioResult field; this field is set to 1 when the call is made, and receives the actual result code upon completion.

As different XPP calls take different arguments in their parameter block, a Pascal variant record has been defined to account for all the different cases. This parameter block is detailed in figure 4. The first four fields (which are the same for all calls) are automatically filled in by the device manager. The CSCode field is automatically filled in by Pascal, depending on which call is being made. The caller must, however, set the ioRefnum field to XPP's reference number, as returned in the OpenXPP call. The ioVRefnum field is unused.

Note that the parameter block is defined so as to be the maximum size used by any call. Different calls take different size parameter blocks, each call requiring a certain minimum size. Callers are free to abbreviate the parameter block where appropriate.

General

With each routine, a list of the parameter block fields used by the call is also given. All routines are invoked by Device Manager Control calls with the csCode field equal to the code corresponding to the function being called. The number next to each field name indicates the byte offset of the field from the start of the parameter block pointed to by A0; only assembly-language programmers need to be concerned with it. . An arrow next to each parameter name indicates whether it's an input, output, or input/output parameter:

Arrow Meaning

- > Parameter is passed
- <-- Parameter is returned
- <--> Parameter is passed and returned

All Device Manager Control calls return an integer result code in the ioResult field. Each routine description lists all the applicable result codes, along with a short description of what the result code means. Refer to the section *XPP Driver result codes* for an alphabetical list of result codes returned by the .XPP driver.

Each routine description includes a Pascal form of the call. Pascal calls to the XPP Driver are of the form:

```
FUNCTION XPPCall (ParamBlock:XPPParamBlkPtr, async:BOOLEAN) : OSErr;
```

XPPCall is the name of the routine.

ParamBlock points to the actual I/O queue element used in the _Control call, filled in by the caller with the parameters of the routine.

async indicates whether or not the call should be made asynchronously. If async is TRUE, the call is executed asynchronously; otherwise the call is executed synchronously.

The routine returns an result code of type OSErr.

```

XPPParamBlock = PACKED RECORD
  qLink:      QElemPtr;      { next queue entry }
  qType:      INTEGER;      { queue type }
  ioTrap:     INTEGER;      { routine trap }
  ioCmdAddr:  Ptr;          { routine address }
  ioCompletion: ProcPtr;    { completion routine }
  ioResult:   OSErr;        { result code }
  CmdResult:  LONGINT;      { Command result (ATP user bytes) [long] }
  ioVRefNum:  INTEGER;      { volume reference or drive number }
  ioRefNum:   INTEGER;      { driver reference number }
  csCode:     INTEGER;      { Call command code }
CASE XPPPrmBlkType OF
  ASPAbortPrm:
    (AbortSCBPtr: Ptr);      { SCB pointer for AbortOS [long] }
  ASPSizeBlk:
    (ASPMaxCmdSize: INTEGER;  { For SPGetParms [word] }
     ASPQuantumSize: INTEGER; { For SPGetParms [word] }
     NumSesss:      INTEGER); { For SPGetParms [word] }
  XPPPrmBlk:
    (SessRefnum:   INTEGER;    { Offset to session refnum [word] }
     ASPTimeout:   Byte;       { Timeout for ATP [byte] }
     ASPRetry:     Byte;       { Retry count for ATP [byte] }
CASE XPPSubPrmType OF
  ASPOpenPrm:
    (ServerAddr: LONGINT;    { Server address block [longword] }
     SCBPointer: Ptr;        { SCB pointer [longword] }
     AttnRoutine: Ptr);      { Attention routine pointer [long] }
  ASPSubPrm:
    (CBSize:      INTEGER;    { Command block size [word] }
     CBPtr:       Ptr;        { Command block pointer [long] }
     RBSize:      INTEGER;    { Reply buffer size [word] }
     RBPtr:       Ptr;        { Reply buffer pointer [long] }
CASE XPPEndPrmType OF
  AFPLoginPrm:
    (AFPAddrBlock: LONGINT; { Address block in AFP login [long] }
     AFPSCBPtr:    Ptr;     { SCB pointer in AFP login [long] }
     AFPAtnRoutine: Ptr);   { Attn routine pointer in AFP login }
  ASPEndPrm:
    (WDSIZE: INTEGER;        { Write data size [word] }
     WDPtr:  Ptr;           { Write data pointer [long] }
     CCBStart: ARRAY[0..295] OF Byte); { CCB memory for driver }
    {Write max size(CCB) = 296; all other calls = 150}
END;

```

Figure 4. .XPP Driver Parameter Block Record

AppleTalk Session Protocol Changes

This section contains descriptions of the XPP driver functions that you can call. Each function description shows the required parameter block fields, their offsets within the parameter block and a brief definition of the field. Possible result codes are also described.

Note on result codes

An important distinction exists between the `ASPParmErr` and `ASPSessClosed` result codes that may be returned by the XPP driver.

When the driver returns `ASPParmErr` to a call that takes as an input a session reference number, the session reference number does not relate to a valid open session. There could be several reasons for this, such as the workstation or server end closed the session or the server end of the session died.

The `ASPSessClosed` result code indicates that even though the session reference number relates to a valid session, that particular session is in the process of closing down (although the session is not yet closed).

OpenSess function

`OpenSess` initiates (opens) a session between the workstation and a server. Below is the required parameter block. A brief definition of the fields follows.

```
FUNCTION ASPOpenSession(xParamBlock:XPParmBlkPtr; async:BOOLEAN): OSErr;
```

Parameter block

--->	26	csCode	word	;always OpenSess
<---	28	SessRefnum	word	;session reference number
--->	30	ASPTimeout	byte	;retry interval in seconds
--->	31	ASPRetry	byte	;number of retries
--->	32	ServerAddr	long word	;server socket address
--->	36	SCBPointer	pointer	;pointer to session control block
--->	40	AttnRoutine	pointer	;pointer to attention routine

`SessRefnum` is a unique number identifying the open session between the workstation and the server. The `SessRefnum` is returned when the function completes successfully and is used in all calls to identify the session.

`ASPTimeOut` is the interval in seconds between retries of the open session request.

`ASPRetry` is the number of retries that will be attempted.

`ServerAddr` is the network identifier or address of the socket on which the server is listening.

`SCBpointer` points to a non-relocatable block of data for the session control block (SCB) that the XPP driver reserves for use in maintaining an open session. The SCB size is defined by the constant `SCBMemSize`. The SCB is a locked block and as long as the session is open, the SCB cannot be modified in any way by the application. There is one SCB for each open session. This block can be reused when a `CloseSess` call is issued and

AppleTalk Manager Update

completed for that session, or when the session is indicated as closed through return of an ASPParamErr as the result of a call for that session.

AttnRoutine is a pointer to a routine that is invoked if an attention from the server is received, or upon session closing. If this pointer is equal to zero, no attention routine will be invoked.

Result codes	ASPNoMoreSess	Driver cannot support another session
	ASPParamErr	Server returned bad (positive) error code
	ASPNoServers	No servers at that address
		The server did not respond to the request.
	ReqAborted	OpenSess was aborted by an AbortOS
	ASPBadvVersNum	Server cannot support the offered version number
	ASPServerBusy	Server cannot open another session

Note: The number of sessions that the driver is capable of supporting depends on the machine that the driver is running on.

CloseSess function

CloseSess closes the session identified by the SessRefnum returned in the OpenSess call. CloseSess aborts any calls that are active on the session, closes the session, and calls the attention routine, if any, with an attention code of zero (zero is invalid as a real attention code).

```
FUNCTION ASPCloseSession(xParamBlock:XPPParamBlkPtr;async:BOOLEAN):  
    OSErr;
```

Parameter block

---	26	csCode	word	;always CloseSess
---	28	SessRefnum	word	;session reference number

Result codes	ASPParamErr	Parameter error, indicates an invalid session reference number
	ASPSessClosed	Session already in process of closing

AbortOS function

AbortOS aborts a pending (not yet completed) OpenSess call. The aborted OpenSess call will return a ReqAborted error.

```
FUNCTION ASPAbortOS(xParamBlock:XPPParamBlkPtr;async:BOOLEAN): OSErr;
```

Parameter block

---	26	csCode	word	;always AbortOS
---	28	AbortSCBPointer	pointer	;pointer to session control block

AbortSCBPointer points to the original SCB used in the the pending OpenSess call.

Result codes `cbNotFound` SCB not found, no outstanding open session to be aborted. Pointer did not point to an open session SCB.

GetParms function

GetParms returns three ASP parameters. This call does not require an open session.

```
FUNCTION ASPGetParms (xParamBlock:XPPParamBlkPtr;async:BOOLEAN): OSErr;
```

Parameter block

---	26	<code>csCode</code>	word	;always GetParms
<---	28	<code>ASPMMaxCmdSize</code>	word	;maximum size of command block
<---	30	<code>ASPQuantumSize</code>	word	;maximum data size
<---	32	<code>NumSesss</code>	word	;number of sessions

`ASPMMaxCmdSize` is the maximum size of a command that can be sent to the server.

`ASPQuantumSize` is the maximum size of data that can be transferred to the server in a Write request or from the server in a command reply.

`NumSesss` is the number of concurrent sessions supported by the driver.

CloseAll function

CloseAll closes every session that the driver has active, aborting all active requests and invoking the attention routines where provided. This call should be used carefully. CloseAll can be used as a system level resource for making sure all sessions are closed prior to closing the driver.

```
FUNCTION ASPCloseAll (xParamBlock:XPPParamBlkPtr;async:BOOLEAN): OSErr;
```

Parameter block

---	26	<code>csCode</code>	word	;always CloseAll
-----	----	---------------------	------	------------------

UserWrite function

UserWrite transfers data on a session. UserWrite is one of the two main calls that can be used to transfer data on an ASP session. The other call that performs a similar data transfer is UserCommand described below. The UserWrite command returns data in two different places. Four bytes of data are returned in the CmdResult field and a variable size reply buffer is also returned.

```
FUNCTION ASPUserWrite (xParamBlock:XPPParamBlkPtr;async:BOOLEAN): OSErr;
```

Parameter block

<---	18	<code>CmdResult</code>	long word	;ASP command result
---	26	<code>csCode</code>	word	;always UserWrite
---	28	<code>SessRefnum</code>	word	;session reference number
---	30	<code>ASPTimeout</code>	byte	;retry interval in seconds

AppleTalk Manager Update

---	32	CBSize	word	;command block size
---	34	CBPtr	pointer	;command block pointer
<-->	38	RBSize	word	;reply buffer size and reply size
---	40	RBPointer	pointer	;reply buffer pointer
<-->	44	WDSize	word	;write data size
---	46	WDPtr	pointer	;write data pointer
<---	50	CCBStart	record	;start of memory for CCB

CmdResult is four bytes of data returned by the server.

SessRefnum is the session reference number returned in the OpenSess call.

ASPTimeOut is the interval in seconds between retries of the call. Notice that there is no ASPRetry field (retries are infinite). The command will be retried at the prescribed interval until completion or the session is closed.

CBSize is the size in bytes of the command data that is to be written on the session. The size of the command block must not exceed the value of ASPMaxCmdSize returned by the GetParms call. Note that this buffer is not the data to be written by the command but only the data of the command itself.

CBPointer points to the command data.

RBSize is passed and indicates the size of the reply buffer in bytes expected by the command. RBSize is also returned and indicates the size of the reply that was actually returned.

RBPointer points to the reply buffer.

WDSize is passed and indicates the size of the write data in bytes to be sent by the command. WDSize is also returned and indicates the size of the write data that was actually written.

WDPointer points to the write data buffer.

CCBStart is the start of the memory to be used by the XPP driver for the command control block. The size of this block is equal to a maximum of 296 bytes. To determine the exact requirement, refer to the CCB Sizes section of this document.

Result codes	ASPParm	Invalid session number, session has been closed
	ASPSizeErr	Command block size is bigger than MaxCmdSize
	ASPSessClosed	Session is closing
	ASPBufTooSmall	Reply is bigger than response buffer
		Buffer will be filled, data will be truncated

UserCommand function

UserCommand is used to send a command to the server on a session.

```
FUNCTION ASPUserCommand(xParamBlock:XPPParamBlkPtr;async:BOOLEAN): OSErr;
```

Parameter block

<---	18	CmdResult	long word	;ASP command result
---	26	csCode	word	;always UserCommand
---	28	SessRefnum	word	;session number
---	30	ASPTimeout	byte	;retry interval in seconds
---	32	CBSize	word	;command block size

--->	34	CBPtr	pointer	;command block pointer
<-->	38	RBSize	word	;reply buffer and reply size
--->	40	RBPointer	pointer	;reply buffer pointer
<---	50	CCBStart	record	;start of memory for CCB

SessRefnum is the session reference number returned in the OpenSess call.

ASPTimeOut is the interval in seconds between retries of the call. Notice that there is no ASPRetry field (retries are infinite). The command will be retried at the prescribed interval until completion or the session is closed.

CBSize is the size in bytes of the block of data that contains the command to be sent to the server on the session. The size of the command block must not exceed the value of ASPMaxCmdSize returned by the GetParms call.

CBPointer points to the block of data containing the command that is to be sent to the server on the session.

RBSize is passed and indicates the size of the reply buffer in bytes expected by the command. RBSize is also returned and indicates the size of the reply that was actually returned.

RBPointer points to the reply buffer.

CCBStart is the start of the memory to be used by the XPP driver for the command control block. The size of this block is equal to a maximum of 150 bytes. To determine the exact requirement refer to the CCB Sizes section of this document.

Result codes	ASPParmErr	Invalid session number, session has been closed
	ASPSizeErr	Command block size is bigger than MaxCmdSize
	ASPSessClosed	Session is closing
	ASPBuTooSmall	Reply is bigger than response buffer
		The buffer will be filled, data will be truncated

GetStatus function

GetStatus returns server status. This call is also used as GetServerInfo at the AFP level. This call is unique in that it transfers data over the network without having a session open. This call does not pass any data but requests that server status be returned.

FUNCTION ASPGetStatus (xParamBlock:XPPParamBlkPtr;async:BOOLEAN): OSErr;

Parameter block

--->	26	csCode	word	;always GetStatus
--->	30	ASPTimeout	byte	;retry interval in seconds
--->	31	ASPRetry	byte	;number of retries
--->	32	ServerAddr	long word	;server socket address
<-->	38	RBSize	word	;reply buffer and reply size
--->	40	RBPointer	pointer	;reply buffer pointer
<---	50	CCBStart	record	;start of memory for CCB

ASPTimeOut is the interval in seconds between retries of the call.

ASPRetry is the number of retries that will be attempted.

AppleTalk Manager Update

ServerAddr is the network identifier or address of the socket on which the server is listening.

RBSize is passed and indicates the size of the reply buffer in bytes expected by the command. RBSize is also returned and indicates the size of the reply that was actually returned.

RBPointer points to the reply buffer.

CCBStart is the start of the memory to be used by the XPP driver for the command control block. The size of this block is equal to a maximum of 200 bytes. To determine the exact requirement refer to the CCB Sizes section of this document.

Result codes	ASPBuTooSmall	Reply is bigger than response buffer Replysize is bigger than ReplyBuffsize.
	ASPNoServer	No response from server at address used in call

AFP implementation

The AFPCall (AFPCommand) function passes a command to an AFP server. The first byte of the AFPCall command buffer (the AFP command byte) must contain a valid AFP command code.

Note: Server information should be gotten through a GetStatus call (described above). GetStatus is equivalent to the AFPGetSrvrInfo. Making an AFP GetSrvrInfo call using AFPCommand, results in a error.

Mapping AFP commands

Most AFP calls are implemented by XPP through a very simple one-to-one mapping of an AFP call to an ASP call without interpretation or verification of the data.

The XPP driver maps AFP command codes to ASP commands according to the following conventions:

AFP Command Code	Comment
\$00	Invalid AFP command
\$01-\$BE (1-190)	Mapped to UserCommand (with the exceptions listed below)
\$BF (191)	Mapped to UserCommand . Reserved for developers; will never be used by Apple
\$C0-\$FD (192-253)	Mapped to UserWrite
\$FE (254)	Mapped to UserWrite; will never be used by Apple
\$FF (255)	Invalid AFP command

The following AFP calls are exceptions to the above conventions:

AFP Command (Code/decimal)	Comment
GetSrvrInfo (15)	Mapped to GetStatus (Use GetStatus to make this call)
Login (18)	Mapped to appropriate login dialog including OpenSess call
LoginCont (19)	Mapped to appropriate login dialog

Logout (20)	Mapped to CloseSess
Write (33)	Mapped to UserWrite

Note that the following AFP calls can pass or return more data than can fit in QuantumSize bytes (8 ATP response packets) and may be broken up by XPP into multiple ASP calls.

AFP Command (Code/decimal)	Comment
Read (27)	Can return up to the number of bytes indicated in ReqCount
Write (33)	Can pass up to the number of bytes indicated in ReqCount

AFPCall function

The AFPCall function can have one of the following command formats.

- General
- Login
- AFPWrite
- AFPRead

General command format

The general command format for the AFPCall function passes an AFP command to the server. This format is used for all AFP calls except AFPLLogin, AFPRead and AFPWrite. Note that from Pascal this call is know as AFPCommand.

```
FUNCTION AFPCommand(xParamBlock:XPPParamBlkPtr;async:BOOLEAN): OSErr;
```

Parameter block

<--	18	CmdResult	long word	;AFP command result
-->	26	csCode	word	;always AFPCall
-->	28	SessRefnum	word	;session reference number
-->	30	ASPTIMEout	byte	;retry interval in seconds
-->	32	CBSize	word	;command buffer size
-->	34	CBPtr	pointer	;command buffer
<-->	38	RBSIZE	word	;reply buffer size and reply size
-->	40	RBPointer	pointer	;reply buffer pointer
<-->	44	WDSIZE	word	;write data size
-->	46	WDPtr	pointer	;write data pointer
<--	50	CCBStart	record	;start of memory for CCB

CmdResult is four bytes of data returned from the server containing an indication of the result of the AFP command.

SessRefnum is the session reference number returned in the AFPLLogin call.

ASPTIMEout is the interval in seconds between retries of the call by the driver.

CBSIZE is the size in bytes of the block of data that contains the command to be sent to the server on the session. The size of the command block must not exceed the value of ASPMaxCmdSize returned by the GetParms call.

AppleTalk Manager Update

CBPointer points to start of the block of data (command block) containing the command that is to be sent to the server on the session. The first byte of the command block must contain the AFP command byte. Subsequent bytes in the command buffer contain the parameters associated with the command as defined in the AFP document.

RBSize is passed and indicates the size of the reply buffer in bytes expected by the command. RBSize is also returned and indicates the size of the reply that was actually returned.

RBPointer points to the reply buffer.

WDSIZE is the size of data to be written to the server (only used if the command is one that is mapped to an ASP UserWrite).

WDPointer points to the write data buffer (only used if the command is one that is mapped to an ASP UserWrite).

CCBStart is the start of the memory to be used by the XPP driver for the command control block. The size of this block is equal to a maximum of 296 bytes. To determine the exact requirement refer to the CCB Sizes section of this document.

Result codes	ASPParmErr	Invalid session number; session has been closed
	ASPSizeErr	Command block size is bigger than MaxCmdSize
	ASPSessClosed	Session is closing
	ASPBufTooSmall	Reply is bigger than response buffer
		Buffer will be filled, data will be truncated
	ParmError	AFP command block size is equal to zero. This error will also be returned if the command byte in the command block is equal to 0 or \$FF (255) or GetSrvrStatus (15).

Login command format

The AFP login command executes a series of AFP operations as defined in the AFP Draft Proposal. For further information, refer to the AFP document.

```
FUNCTION AFPCmd(xParamBlock:XPPParamBlkPtr;async:BOOLEAN) : OSErr;
```

Parameter block

<--	18	CmdResult	long word	;AFP command result
-->	26	csCode	word	;always AFPCall
<--	28	SessRefnum	word	;session reference number
-->	30	ASPTimeout	byte	;retry interval in seconds
-->	31	ASPRetry	byte	;number of retries
-->	32	CBSize	word	;command buffer size
-->	34	CBPtr	pointer	;command buffer
<-->	38	RBSize	word	;reply buffer size and reply size
-->	40	RBPtr	pointer	;reply buffer pointer
-->	44	AFPAddrBlock	long word	;server address block
<-->	48	AFPSCBPointer	pointer	;SCB pointer
<-->	52	AFPAttnRoutine	pointer	;attention routine pointer
<--	50	CCBStart	record	;start of command control block

CmdResult is four bytes of data returned from the server containing an indication of the result of the AFP command.

SessRefnum is the session reference number (returned by the AFPLLogin call).

ASPTimeOut is the interval in seconds between retries of the call.

ASPRetry is the number of retries that will be attempted.

CBSize is the size in bytes of the block data that contains the command to be sent to the server on the session. The size of the command block must not exceed the value of ASPMaxCmdSize returned by the GetParms call.

CBPointer points to the block of data (command block) containing the AFP login command that is to be sent to the server on the session. The first byte of the command block must be the AFP login command byte. Subsequent bytes in the command buffer contain the parameters associated with the command.

RBSize is passed and indicates the size of the reply buffer in bytes expected by the command. RBSize is also returned and indicates the size of the reply that was actually returned.

RBPointer points to the reply buffer.

AFPServerAddr is the network identifier or address of the socket on which the server is listening.

AFPSCBPointer points to a locked block of data for the session control block (SCB). The SCB size is defined by SCBMemSize. The SCB is a locked block and as long as the session is open, the SCB can not be modified in any way by the application. There is one SCB for each open session.

AFPAttnRoutine is a pointer to a routine that is invoked if an attention from the server is received. When AFPAttnRoutine is equal to zero, no attention routine will be invoked.

CCBStart is the start of the memory to be used by the XPP driver for the command control block. The size of this block is equal to a maximum of 150 bytes. To determine the exact requirement refer to the CCB Sizes section of this document.

Note : In the parameter block, the AFPSCBPointer and the AFPAttnRoutine fields overlap with the start of the CCB and are modified by the call.

Result codes	ASPSizeErr	Command block size is bigger than MaxCmdSize
	ASPBufTooSmall	Reply is bigger than response buffer
		Buffer will be filled, data will be truncated
	ASPNoServer	Server not responding
	ASPServerBusy	Server cannot open another session
	ASPSBadVersNum	Server cannot support the offered ASP version number
	ASPNoMoreSess	Driver cannot support another session.

AFPWrite command format

The AFPWrite and AFPRead command formats allow the calling application to make AFP-level calls that read or write a data block that is larger than a single ASP-level call is capable of reading or writing. The maximum number of bytes of data that can be read or written at the ASP level is equal to QuantumSize.

```
FUNCTION AFPCommand(xParamBlock:XPPParamBlkPtr;async:BOOLEAN): OSErr;
```

Parameter block

AppleTalk Manager Update

<---	18	CmdResult	long word	;AFP command result
---	26	csCode	word	;always AFPCall
---	28	SessRefnum	word	;session number
---	30	ASPTIMEout	byte	;retry interval in seconds
---	32	CBSize	word	;command buffer size
---	34	CBPtr	pointer	;command buffer
<-->	38	RBSize	word	;reply buffer size and reply size
---	40	RBPtr	pointer	;reply buffer pointer
<---	44	WDSIZE	word	;(used internally)
<-->	46	WDPtr	pointer	;write data pointer (updated)
<---	50	CCBStart	record	;start of memory for CCB

CmdResult is four bytes of data returned from the server containing an indication of the result of the AFP command.

SessRefnum is the session reference number returned in the AFPLLogin call.

ASPTIMEout is the interval in seconds between retries of the call.

CBSize is the size in bytes of the block data that contains the command to be sent to the server on the session. The size of the command block must not exceed the value of ASPMaxCmdSize returned by the ASPGetParms call.

CBPointer points to the block of data (see command block structure below) containing the AFP write command that is to be sent to the server on the session. The first byte of the Command Block must contain the AFP write command byte.

RBSize is passed and indicates the size of the reply buffer in bytes expected by the command. RBSize is also returned and indicates the size of the reply that was actually returned.

RBPter points to the reply buffer.

WDSIZE is used internally.

Note: This command does not pass the write data size in the queue element but in the command buffer. XPP will look for the size in that buffer.

WDPtr is a pointer to the block of data to be written. Note that this field will be updated by XPP as it proceeds and will always point to that section of the data which XPP is currently writing.

CCBStart is the start of the memory to be used by the XPP driver for the command control block. The size of this block is equal to a maximum of 296 bytes. To determine the exact requirement refer to the CCB Sizes section of this document.

Command block structure

The AFP write command passes several arguments to XPP in the command buffer itself. The byte offsets are relative to the location pointed to by CBPointer.

---	0	CmdByte	byte	;AFP call command byte
---	1	StartEndFlag	byte	;Start/end Flag
<-->	4	RWOffset	long word	;offset within fork to write
<-->	8	ReqCount	long word	;requested count

CmdByte is the AFP call command byte and must contain the AFP write command code.

StartEndFlag is a one-bit flag (the high bit of the byte) indicating whether the RWOOffset field is relative to the beginning or the end of the fork (all other bits are zero).

0 = relative to the beginning of the fork

1 = relative to the end of the fork

RWOOffset is the byte offset within the fork at which the write is to begin.

ReqCount indicates the size of the data to be written and is returned as the actual size written.

Note that the RWOOffset and ReqCount fields are modified by XPP as the write proceeds and will always indicate the current value of these fields.

The Pascal structure of the AFP command buffer follows:

```
AFPCommandBlock = PACKED RECORD
    CmdByte:      Byte;
    StartEndFlag: Byte;
    ForkRefNum:   INTEGER;    {Used by server}
    RWOOffset:    LONGINT;
    ReqCount:     LONGINT;
    NewLineFlag:  Byte;       {Unused by write}
    NewLineChar:  CHAR;       {Unused by write}
END;
```

Result codes	ASPParamErr	Invalid session number
	ASPSizeErr	Command block size is bigger than MaxCmdSize
	ASPSessClosed	Session is closing
	ASPBuTooSmall	Reply is bigger than response buffer

AFPRead command format

```
FUNCTION AFPCommand(xParamBlock:XPPParamBlkPtr;async:BOOLEAN): OSErr;
```

Parameter block

<---	18	CmdResult	long word	;ASP command result
---	26	csCode	word	;always AFPCall
---	28	SessRefnum	word	;session number
---	30	ASPTIMEOUT	byte	;retry interval in seconds
---	32	CBSIZE	word	;command buffer size
---	34	CBPTR	pointer	;command buffer
---	38	RBSIZE	word	;used internally
<-->	40	RBPTR	pointer	;reply buffer pointer (updated)
<---	50	CCBSTART	record	;start of memory for CCB

CmdResult is four bytes of data returned from the server containing an indication of the result of the AFP command.

SessRefnum is the session reference number returned in the AFPLLogin call.

ASPTIMEOUT is the interval in seconds between retries of the call.

CBSize is the size in bytes of the block data that contains the command to be sent to the server on the session. The size of the command block must not exceed the value of ASPMaxCmdSize returned by the GetParms call.

CBPointer points to the block of data (command block) containing the AFP read command that is to be sent to the server on the session. The first byte of the command block must contain the AFP read command byte. The command block structure is shown below.

RBSize is used internally.

Note: This command does not pass the read size in the queue element but in the command buffer. XPP will look for the size in that buffer.

RBPointer points to the reply buffer. Note that this field will be updated by XPP as it proceeds and will always point to that part of the buffer that XPP is currently reading into.

CCBStart is the start of the memory to be used by the XPP driver for the command control block. The size of this block is equal to a maximum of 150 bytes. To determine the exact requirement refer to The CCB Sizes section .

Command block structure

The AFP read command passes several arguments to XPP in the command buffer itself. The byte offsets are relative to the location pointed to by CBPointer.

---	0	CmdByte	byte	;AFP call command byte
<-->	4	RWOffset	long word	;offset within fork to read
<-->	8	ReqCount	long word	;requested count
---	12	NewLineFlag	byte	;Newline Flag
---	13	NewLineChar	byte	;Newline Character

CmdByte is the AFP call command byte and must contain the AFP read command code.

RWOffset is the byte offset within the fork at which the read is to begin.

ReqCount indicates the size of the read data buffer and is returned as the actual size read.

Note that the RWOffset and ReqCount fields are modified by XPP as the read proceeds and will always indicate the current value of these fields.

NewLineFlag is a one-bit flag (the high bit of the byte) indicating whether or not the read is to terminate at a specified character (all other bits are zero).

0 = no Newline Character is specified

1 = a Newline Character is specified

NewLineChar is any character from \$00 to \$FF (inclusive) which when encountered in reading the fork, causes the read operation to terminate.

The Pascal structure of the AFP command buffer follows:

```
AFPCommandBlock = PACKED RECORD
    CmdByte:      Byte;
    StartEndFlag: Byte;      {Unused for read}
    ForkRefNum:   INTEGER;   {Used by server}
    RWOffset:     LONGINT;
    ReqCount:     LONGINT;
    NewLineFlag:  Byte;
    NewLineChar:  CHAR;
END;
```


Result codes	ASPPParamErr	Invalid session number
	ASPSizeErr	Command block size is bigger than MaxCmdSize
	ASPSessClosed	Session is closing
	ASPBufTooSmall	Reply is bigger than response buffer

CCB sizes

The XPP driver uses the memory provided at the end of the UserWrite, UserCommand and GetStatus parameter blocks as an internal command control block (CCB). Using the maximum block sizes specified in the call descriptions will provide adequate space for the call to execute successfully. However, this section is provided for developers who wish to minimize the amount of memory taken up by the CCB in the queue element.

Specifically, this memory is used for creating data structures to be used in making calls to the ATP driver. This includes parameter blocks and buffer data structures (BDS's – detailed in the AppleTalk Manager chapter of *Inside Macintosh*). The exact size of this memory depends on the size of the response expected, and, in the case of UserWrite, on the size of data to be written.

In the UserCommand and GetStatus cases (along with all AFP calls which map to UserCommand), a BDS must be set up to hold the response information. The number of entries in this BDS is equal to the size of the response buffer divided by the maximum number of data bytes per ATP response packet (578), rounded up. Note that as described in the ASP document in *Inside AppleTalk*, ASP must ask for an extra response in the case where the response buffer is an exact multiple of 578. Of course, no BDS can be more than eight elements big. XPP also needs bytes for the queue element to call ATP with, so the minimum size of a CCB, as a function of the response buffer size (RBSIZE) is:

$$\begin{aligned} \text{BDSSize} &= \text{MIN}(((\text{RBSIZE DIV } 578) + 1), 8) * \text{BDSEntrySz} \\ \text{CCBSIZE} &= \text{IOQEISize} + 4 + \text{BDSSize} \end{aligned}$$

In the UserWrite (and AFP calls mapping to UserWrite) case, XPP needs to create an additional BDS and queue element to use in sending the write data to the server. Thus the minimum size of a UserWrite CCB, as a function of the response buffer and write data sizes (RBSIZE and WDSIZE) is:

$$\begin{aligned} \text{WrBDSSize} &= \text{MIN}(((\text{WDSIZE DIV } 578) + 1), 8) * \text{BDSEntrySz} \\ \text{WrCCBSIZE} &= (2 * \text{IOQEISize}) + 4 + \text{BDSSize} + \text{WrBDSSize} \end{aligned}$$

Note: BDSEntrySz is equal to 12. IOQEISize is equal to 50.

AFP Driver Result Codes

Result Code	Comment	Returned by
ASPSBadVersNum	Server cannot support the offered version number	OpenSess AFPCall (Login)
ASPSBufTooSmall	Reply is bigger than response buffer Buffer will be filled, data may be truncated.	UserWrite UserCommand GetStatus AFPCall
ASPSNoMoreSess	Driver cannot support another session	OpenSess AFPCall (Login)
ASPSNoServers	No servers at that address The server did not respond to the request.	GetStatus OpenSess AFPCall (Login)
ASPSParamErr	Parameter error server returned bad (positive) error code Invalid Session Reference Number	OpenSess CloseSess UserWrite UserCommand AFPCall
ASPSServerBusy	Server cannot open another session	OpenSess AFPCall (Login)
ASPSsessClosed	Session already in process of closing	CloseSess UserWrite UserCommand AFPCall
ASPSsizeErr	Command block size is bigger than MaxParamSize	UserWrite UserCommand AFPCall
cbNotFound	SCB not found, no outstanding open session to be aborted. Pointer did not point to an open session SCB.	AbortOS
ParmError	AFP Command Block size is less than or equal to zero. Command byte in the Command block is equal to 0 or \$FF (255) or GetSrvrStatus (15).	AFPCall
ReqAborted	Open session was aborted by an Abort Open Session	OpenSess AFPCall (Login)

SUMMARY

SetSelfSend function

Parameter Block

--->	26	csCode	word	; always SetSelfSend
--->	28	NewSelfFlag	byte	; new SelfSend flag
<---	29	OldSelfFlag	byte	; old SelfSend flag

AppleTalk Transaction Protocol

NSendRequest function

Parameter block

--->	18	UserData	longword	;user bytes
<---	22	reqTID	word	;transaction ID used in request
--->	26	csCode	word	;always sendRequest
--->	28	atpSocket	byte	;socket to send request on
<-->	29	atpFlags	byte	;control information
--->	30	addrBlock	longword	;destination socket address
--->	34	reqLength	word	;request size in bytes
--->	36	reqPointer	pointer	;pointer to request data
--->	40	bdsPointer	pointer	;pointer to response BDS
--->	44	numOfBufs	byte	;number of responses expected
--->	45	timeOutVal	byte	;timeout interval
<---	46	numOf Resps	byte	;number of responses received
<-->	47	retryCount	byte	;number of retries
<---	48	intBuff	word	;used internally

KillSendReq function

Parameter block

--->	26	csCode	word	; always KillSendReq
--->	44	AKillQEI	pointer	; pointer to queue element

KillGetReq function

Parameter block

--->	26	csCode	word	; always KillGetReq
--->	44	AKillQEI	pointer	; pointer to queue element

Name Binding Protocol

KillNBP function

Parameter block

---	26	csCode	word	; always KillNBP
---	28	NKillQE1	pointer	; pointer to queue element

Data Types

XPPParamBlock = PACKED RECORD

qLink:	QElemPtr;	{ next queue entry }
qType:	INTEGER;	{ queue type }
ioTrap:	INTEGER;	{ routine trap }
ioCmdAddr:	Ptr;	{ routine address }
ioCompletion:	ProcPtr;	{ completion routine }
ioResult:	OSErr;	{ result code }
CmdResult:	LONGINT;	{ Command result(ATP user bytes)[long]}
ioVRefNum:	INTEGER;	{ volume reference or drive number }
ioRefNum:	INTEGER;	{ driver reference number }
csCode:	INTEGER;	{ Call command code }
CASE XPPPrmBlkType OF		
ASPAbsortPrm:		
(AbsortSCBPtr: Ptr);		{ SCB pointer for AbortOS [long] }
ASPSizeBlk:		
(ASPMaxCmdSize: INTEGER;		{ For SPGetParms [word] }
ASPQuantumSize: INTEGER;		{ For SPGetParms [word] }
NumSesss: INTEGER);		{ For SPGetParms [word] }
XPPPrmBlk:		
(SessRefnum: INTEGER;		{ Offset to session refnum [word] }
ASPTimeout: Byte;		{ Timeout for ATP [byte] }
ASPRetry: Byte;		{ Retry count for ATP [byte] }
CASE XPPSubPrmType OF		
ASPOpenPrm:		
(ServerAddr: LONGINT;		{ Server address block [longword] }
SCBPointer: Ptr;		{ SCB pointer [longword] }
AttnRoutine: Ptr);		{ Attention routine pointer [long] }
ASPSubPrm:		
(CBSize: INTEGER;		{ Command block size [word] }
CBPtr: Ptr;		{ Command block pointer [long] }
RBSize: INTEGER;		{ Reply buffer size [word] }
RBPtr: Ptr;		{ Reply buffer pointer [long] }
CASE XPPEndPrmType OF		
AFPLoginPrm:		
(AFPAAddrBlock: LONGINT;		{ Address block in AFP login [long] }
AFPSCBPtr: Ptr;		{ SCB pointer in AFP login [long] }
AFPAttnRoutine: Ptr);		{ Attn routine pointer in AFP login }
ASPEndPrm:		
(WDSize: INTEGER;		{ Write data size [word] }
WDPtr: Ptr;		{ Write data pointer [long] }
CCBStart: ARRAY[0..295] OF Byte);		{ CCB memory for driver }

END;

```

AFPCCommandBlock = PACKED RECORD
  CmdByte:      Byte;
  StartEndFlag: Byte;      {Unused for read}
  ForkRefNum:   INTEGER;   {Used by server}
  RWOOffset:    LONGINT;
  ReqCount:     LONGINT;
  NewLineFlag:  Byte;      {Unused by write}
  NewLineChar:  CHAR;      {Unused by write}
END;

```

AppleTalk Session Protocol

```

FUNCTION ASPOpenSession(xParamBlock:XPPParamBlkPtr;async:BOOLEAN) OSErr;

```

Parameter block

--->	26	csCode	word	;always OpenSess
<---	28	SessRefnum	word	;session reference number
--->	30	ASPTimeout	byte	;retry interval in seconds
--->	31	ASPRetry	byte	;number of retries
--->	32	ServerAddr	long word	;server socket address (SLS)
--->	36	SCBPointer	pointer	;pointer to session control block
--->	40	AttnRoutine	pointer	;pointer to attention routine

```

FUNCTION ASPCloseSession(xParamBlock:XPPParamBlkPtr;async:BOOLEAN) OSErr;

```

Parameter block

--->	26	csCode	word	;always CloseSess
--->	28	SessRefnum	word	;session reference number

```

FUNCTION ASPAbortOS(xParamBlock:XPPParamBlkPtr;async:BOOLEAN) OSErr;

```

Parameter block

--->	26	csCode	word	;always AbortOS
--->	28	AbortSCBPtr	pointer	;pointer to session control block

```

FUNCTION ASPGetParms(xParamBlock:XPPParamBlkPtr;async:BOOLEAN) OSErr;

```

Parameter block

--->	26	csCode	word	;always GetParms
<---	28	ASPMaxCmdSize	word	;maximum size of command block
<---	30	ASPQuantumSize	word	;maximum data size
<---	32	NumSesss	word	;number of sessions

```

FUNCTION ASPCloseAll(xParamBlock:XPPParamBlkPtr;async:BOOLEAN) OSErr;

```

Parameter block

--->	26	csCode	word	;always CloseAll
------	----	--------	------	------------------

AppleTalk Manager Update

FUNCTION ASPUserWrite(xParamBlock:XPPParmBlkPtr;async:BOOLEAN) OSErr;

Parameter block

<---	18	CmdResult	long word	;ASP command result
---	26	csCode	word	;always UserWrite
---	28	SessRefnum	word	;session reference number
---	30	ASPTimeout	byte	;retry interval in seconds
---	32	CBSize	word	;command block size
---	34	CBPtr	pointer	;command block pointer
<-->	38	RBSize	word	;reply buffer size and reply size
---	40	RBPtr	pointer	;reply buffer pointer
<-->	44	WDSIZE	word	;write data size
---	46	WDPtr	pointer	;write data pointer
<---	50	CCBStart	record	;start of memory for CCB

FUNCTION ASPUserCommand(xParamBlock:XPPParmBlkPtr;async:BOOLEAN) OSErr;

Parameter block

<---	18	CmdResult	long word	;ASP command result
---	26	csCode	word	;always UserCommand
---	28	SessRefnum	word	;session number
---	30	ASPTimeout	byte	;retry interval in seconds
---	32	CBSize	word	;command block size
---	34	CBPtr	pointer	;command block pointer
<-->	38	RBSize	word	;reply buffer and reply size
---	40	RBPtr	pointer	;reply buffer pointer
<---	50	CCBStart	record	;start of memory for CCB

FUNCTION ASPGetStatus(xParamBlock:XPPParmBlkPtr;async:BOOLEAN) OSErr;

Parameter block

---	26	csCode	word	;always GetStatus
---	30	ASPTimeout	byte	;retry interval in seconds
---	31	ASPRetry	byte	;number of retries
---	32	ServerAddr	long word	;server socket address
<-->	38	RBSize	word	;reply buffer and reply size
---	40	RBPtr	pointer	;reply buffer pointer
<---	50	CCBStart	record	;start of memory for CCB

AppleTalk Filing Protocol

AFPCall Function

General command format

```
FUNCTION AFPCall(xParamBlock:XPPParamBlkPtr;async:BOOLEAN) OSErr;
```

Parameter block

<---	18	CmdResult	long word	;AFP command result
--->	26	csCode	word	;always AFPCall
--->	28	SessRefnum	word	;session reference number
--->	30	ASPTIMEOUT	byte	;retry interval in seconds
--->	32	CBSIZE	word	;command buffer size
--->	34	CBPtr	pointer	;command buffer
<-->	38	RBSIZE	word	;reply buffer size and reply size
--->	40	RBPTR	pointer	;reply buffer pointer
<-->	44	WDSIZE	word	;write data size
--->	46	WDPtr	pointer	;write data pointer
<---	50	CCBStart	record	;start of memory for CCB

Login command format

```
FUNCTION AFPCall(xParamBlock:XPPParamBlkPtr;async:BOOLEAN) OSErr;
```

Parameter block

<---	18	CmdResult	long word	;AFP command result
--->	26	csCode	word	;always AFPCall
<---	28	SessRefnum	word	;session reference number
--->	30	ASPTIMEOUT	byte	;retry interval in seconds
--->	31	ASPRetry	byte	;number of retries
--->	32	CBSIZE	word	;command buffer size
--->	34	CBPtr	pointer	;command buffer
<-->	38	RBSIZE	word	;reply buffer size and reply size
--->	40	RBPTR	pointer	;reply buffer pointer
--->	44	AFPAAddrBlock	long word	;server address block
<-->	48	AFPSCBPointer	pointer	;SCB pointer
<-->	52	AFPAtnRoutine	pointer	;attention routine pointer
<---	50	CCBStart	record	;start of command control block

AFPWrite command format

```
FUNCTION AFPCall(xParamBlock:XPPParamBlkPtr;async:BOOLEAN) OSErr;
```

Parameter block

<---	18	CmdResult	long word	;AFP command result
--->	26	csCode	word	;always AFPCall
--->	28	SessRefnum	word	;session number
--->	30	ASPTIMEOUT	byte	;retry interval in seconds
--->	32	CBSIZE	word	;command buffer size
--->	34	CBPtr	pointer	;command buffer

AppleTalk Manager Update

<-->	38	RBSize	word	;reply buffer size and reply size
-->	40	RBPtr	pointer	;reply buffer pointer
<--	44	WDSIZE	word	; (used internally)
<-->	46	WDPtr	pointer	;write data pointer (updated)
<--	50	CCBStart	record	;start of memory for CCB

Command block structure

-->	0	CmdByte	byte	;AFP call command byte
-->	1	StartEndFlag	byte	;Start/end Flag
<-->	4	RWOffset	long word	;offset within fork to write
<-->	8	ReqCount	long word	;requested count

AFPRead command format

FUNCTION AFPCCommand(xParamBlock:XPPParamBlkPtr; async:BOOLEAN) OSerr;

Parameter block

<--	18	CmdResult	long word	;ASP command result
-->	26	csCode	word	;always AFPCall
-->	28	SessRefnum	word	;session number
-->	30	ASPTIMEOUT	byte	;retry interval in seconds
-->	32	CBSIZE	word	;command buffer size
-->	34	CBPtr	pointer	;command buffer
-->	38	RBSize	word	;used internally
<-->	40	RBPtr	pointer	;reply buffer pointer (updated)
<--	50	CCBStart	record	;start of memory for CCB

Command block structure

The AFP read command passes several arguments to XPP in the command buffer itself. The byte offsets are relative to the location pointed to by *CBPtr*.

-->	0	CmdByte	byte	;AFP call command byte
<-->	4	RWOffset	long word	;offset within fork to read
<-->	8	ReqCount	long word	;requested count
-->	12	NewLineFlag	byte	;Newline Flag
-->	13	NewLineChar	byte	;Newline Character

AppleTalk Session Protocol Constants

Offsets in user bytes

ASPCmdCode	EQU	0	; Offset to command field
ASPWSSNum	EQU	1	; WSS number in OpenSessions
ASPVersNum	EQU	2	; ASP version number in OpenSessions
ASPSSSNum	EQU	0	; SSS number in OpenSessReplies
ASPSessID	EQU	1	; Session ID (requests &OpenSessReply)
ASPOpenErr	EQU	2	; OpenSessReply error code
ASPSeqNum	EQU	2	; Sequence number in requests

ASPAttnCode EQU 2 ; Attention bytes in attentions

Offsets in ATP data part

ASPWrBSize EQU 0 ; Offset to write buffer size (WriteData)
 ASPWrHdrSz EQU ASPWrBSize+2 ; Size of data part

ASP command codes

ASPCloseSess EQU 1 ; Close session
 ASPCommand EQU 2 ; User-command
 ASPGetStat EQU 3 ; Get status
 ASPOpenSess EQU 4 ; Open session
 ASPTickle EQU 5 ; Tickle
 ASPWrite EQU 6 ; Write
 ASPDataWrite EQU 7 ; WriteData (from server)
 ASPAttention EQU 8 ; Attention (from server)

ASP miscellaneous

ASPVersion EQU \$0100 ; ASP version number
 MaxCmdSize EQU ATPMaxData ; Maximum command block size
 QuantumSize EQU ATPMaxData*ATPMaxNum ; Maximum reply size
 XPPLoadedBit EQU ATPLoadedBit+1 ; XPP bit in PortBUse
 XPPUnitNum EQU 40 ; Unit number for XPP (old ROMs)

ASP errors codes

ASPSBadVersNum EQU -1066 ; Server cannot support this ASP version
 ASPBufTooSmall EQU -1067 ; Buffer too small
 ASPNoMoreSess EQU -1068 ; No more sessions on server
 ASPNoServers EQU -1069 ; No servers at that address
 ASPParamErr EQU -1070 ; Parameter error
 ASPServerBusy EQU -1071 ; Server cannot open another session
 ASPSessClosed EQU -1072 ; Session closed
 ASPSizeErr EQU -1073 ; Command block too big

Control codes

OpenSess EQU 255 ; Open session
 CloseSess EQU 254 ; Close session
 UserCommand EQU 253 ; User command
 UserWrite EQU 252 ; User write
 GetStatus EQU 251 ; Get status
 AFPCall EQU 250 ; AFP command (buffer has command code)
 GetParms EQU 249 ; Get parameters
 AbortOS EQU 248 ; Abort open session request
 CloseAll EQU 247 ; Close all open sessions

ASP queue element standard structure: arguments passed in the CSPParam area

SessRefnum EQU CParam ; Offset to session refnum [word]
 ASPTimeout EQU SessRefnum+2 ; Timeout for ATP [byte]
 ASPRetry EQU ASPTimeout+1 ; Retry count for ATP [byte]
 ServerAddr EQU ASPRetry+1 ; Server address block [longword]

AppleTalk Manager Update

SCBPointer	EQU	ServerAddr+4 ; SCB pointer [longword]
AttnRoutine	EQU	SCBPointer+4 ; Attention routine pointer [long]
CBSize	EQU	ASPRetry+1 ; Command block size [word]
CBPtr	EQU	CBSize+2 ; Command block pointer [long]
RBSize	EQU	CBPtr+4 ; Reply buffer size [word]
RBPtr	EQU	RBSize+2 ; Reply buffer pointer [long]
WDSize	EQU	RBPtr+4 ; Write data size [word]
WDPtr	EQU	WDSize+2 ; Write data pointer [long]
CCBStart	EQU	WDPtr+4 ; Start of memory for CCB
ASPMxCmdSize	EQU	CSPParam ; For SPGetParms [word]
ASPQuantumSize	EQU	ASPMxCmdSize+2 ; For SPGetParms [word]
AbortSCBPtr	EQU	CSPParam ; SCB pointer for AbortOS [long]
CmdResult	EQU	UserData ; Command result (ATP user bytes) [long]
AFPAddrBlock	EQU	RBPtr+4 ; Address block in AFP login [long]
AFPSCBPtr	EQU	AFPAddrBlock+4 ; SCB pointer in AFP login [long]
AFPAttnRoutine	EQU	AFPSCBPtr+4 ; Attn routine pointer in AFP login
SCBMemSize	EQU	\$C0 ; Size of memory for SCB

AFPCall command codes

AFPLogin	EQU	18;
AFPContLogin	EQU	19;
AFPLogout	EQU	20;
AFPRead	EQU	27;
AFPWrite	EQU	33;

Offsets for certain parameters in Read/Write calls

StartEndFlag	EQU	1 ; Write only; offset relative to start or end
RWOffset	EQU	4 ; Offset at which to start read or write
ReqCount	EQU	8 ; Count of bytes to read or write
NewLineFlag	EQU	12 ; Read only; newline character flag
NewLineChar	EQU	13 ; Read only; newline character
LastWritten	EQU	0 ; Write only; last written (returned)

Miscellaneous

AFPUseWrite	EQU	\$C0; first call in range that maps to an ASPWrite
-------------	-----	--

